Preface

This book teaches introductory CMOS digital electronics for electrical and computer engineering undergraduates. For many years the CMOS technology dominated the method of designing and manufacturing digital (computing) integrated circuits. The selection of material here is not significantly different than the graduate texts by J. Rabaey et al, N. Weste and D. Harris, J. Ureyama, or J. Baker, but the style is introductory with many examples, self-exercises, and end-of-chapter problems.

This book initially reviews material relevant to digital electronics that students learned in previous circuit and logic courses. The book then moves through chapters on: basic physics of semiconductor materials and diodes; *n*MOS and *p*MOS field effect transistors circuit analysis; electronic properties of the metal interconnections; the CMOS inverter; the CMOS NAND, NOR, and transmission gates electronics; transformation from Boolean equations to CMOS transistor schematics and domino circuits; timing electronics; memory circuits; FPGAs; CMOS layout; and CMOS fabrication basics. The emphasis is on transistor level electronics.

The principles of power dissipation are introduced with numerical examples. Lowering circuit power has special urgency today where total Internet power consumes about 10% of USA electrical power generation.

Other features and objectives include

- Abundant examples, self-exercises with answers, and many problems at the end of chapters to give students reflexive skills in transistor circuit analysis
- The material can be used before or after a companion class in introductory analog electronics
- The book strives for clarity and self-learning in an undergraduate presentation
- Don’t overwhelm students with too much detail; define teaching goals consistent with what they will take forward to the next level of electronics
- Provide students with an education that serves as a prerequisite for graduate or senior courses in digital electronics and allows entry level into the digital electronics industry

- The book is light enough for students to carry to class

A summary description of each chapter is

Chapter 1 Basic Logic Gate and Circuit Theory

The chapter reviews relevant logic theory that includes Boolean equation to logic gate schematics, DeMorgan's Theorem, logic equivalence, and logic gate reduction. Basic circuit theory is next including resistor power supply circuits with emphasis on analysis of terminal impedance, node voltages, and branch currents by inspection. Non-linear circuit analysis techniques are introduced using the diode and its non-linear current-voltage expression. Capacitor and inductor properties and circuits are analyzed, as is the power wasted in resistive, capacitive, and inductive circuits. These review topics are a few among many, but are selected for their relevance in the digital circuit analysis that follows.

Chapter 2 Semiconductor Physics

Chapter 2 introduces the semiconductor physics that underlies device operation. The goal is to impart a good visual model of the physics of materials and diodes, and to use basic equations for better understanding. Semiconductor physics is a complex subject that can involve more than one course at the graduate level, and Chapter 2 cannot replicate this. But visual models of semiconductor materials and diodes are important since engineers often use qualitative language to communicate important properties of the physics of semiconductor diodes and transistors. Students should be able to answer the question, "How do diodes (and transistors) work?" and perform basic parameter calculations. This chapter leads directly into Chapter 3 on field effect transistors.

Chapter 3 MOSFET Transistors

CMOS circuits use two forms of transistors; the *n*MOS and *p*MOS field effect transistors. Chapter 3 describes how these transistors work followed by numerical analysis of circuit node voltages and currents. Many examples, self-exercises, and end-of-chapter problems give students the reflexive response to analyze transistor digital circuits. Equal treatment is given to each transistor type.

Chapter 4 IC Metal Interconnection Properties

Metal properties are especially relevant in modern circuits since chip total metal length may be on the order of several miles, and minimum metal dimensions can be 45 nm or smaller.

Metal properties are a dominant factor in attaining maximum IC frequency and minimum noise operation, and metal physics deserves as much study as does the transistor.

Chapter 5 The CMOS Inverter

The CMOS inverter is the most abundant logic gate in any digital IC. It has one *n*MOS and one *p*MOS field effect transistor. It has about a dozen important electronic properties that are introduced with numerical examples. Inverter properties are inherently important, but are also the basis for electronic properties of NAND, NOR, and sequential logic circuits such as the master-slave flip-flop. Inverter power dissipation properties are emphasized.

Chapter 6 The CMOS NAND, NOR, and Transmission Gates

NAND and NOR gates build on the inverter by adding transistors in parallel and series to the inverter. These multi-input logic gates have all of the electronic properties of the inverter and a few that are unique. The electronic basis for the non-controlling logic state is described as it relates to circuit debugging, test engineering, and schematic reading. Pass transistor and CMOS transmission gate properties conclude the chapter. Transmission gates are abundant comprising half of the logic gates in the master-slave flip-flop.

Chapter 7 CMOS Design Styles

Chapter 7 develops design styles that assemble transistors into logic gates. It begins with a relatively simple technique to transform Boolean equations into a transistor schematic that performs the logic. Other design styles are presented along with the reasons for having different styles. Power dissipation is analyzed with a technique that allows a power comparison of different combinational logic configurations.

Chapter 8 Sequential Gate Design and Timing

The accurate design and placement of timing signals may be the most challenging task for a designer. This often neglected undergraduate course topic is emphasized giving it the importance it deserves. The edge-triggered FF has a complexity that must be mastered. Timing parameters and rules must be exact or the circuit will fail. System level timing builds on these foundations and introduces system timing parameters and constraints.

Chapter 9 Memory Circuits

Memory circuits have always been embedded within the computing chips. Today microprocessor chips may dedicate more than 65% of the total transistor count to these memory circuits. Therefore special emphasis is placed on these static random access memory

(SRAM) designs. Transistor sizing of SRAM cells is developed with numerical examples. Another high volume memory design is the dynamic random access memory called the DRAM. This single transistor memory cell has different properties.

Chapter 10 Field Programmable Arrays (FPGAs)

Chapter 8 looks at a unique and popular design style using Field Programmable Gate Arrays (FPGA). This material follows from other design styles described in the preceding chapter. The electronics and method of operation are different, but FPGAs are common and abundant enough to devote a chapter.

Chapter 11 CMOS Layout

A conversion occurs in the design process when transistor schematics are transformed to rectangular images on a photographic mask. The images represent transistor and metal line geometries. Masks are drawn for each of several layers in the buildup of the IC. Layout is not electronics, but is the necessary first step in using photolithography to make the tiny transistors and metal interconnections. The mask layout step is introduced using manual layout of the inverter, NAND, and NOR gates. Several commercial layout tools exist, but cost and training time led us to consider the Microsoft PowerPoint program to draw the layouts. PowerPoint is typically available on all computers, training time is minimal, it appears to have long-term stability in the market, and students get a better grounding in design rules. PowerPoint has been successful as a teaching tool in the classroom for layout of simple logic gates circuits.

Chapter 12 How Chips are Made

Chapter 12 describes the chemical, physical, and photolithography techniques that actually make the final circuit. This chapter is qualitative, but sufficient to allow students to converse on the various sequenced fabrication techniques that achieve the end circuit result of the chip.

Comments for Instructors

The book uses long channel models for MOSFET analysis even though short channel models are common in industry. The reasons are twofold. First, the short channel models are often simplified for undergraduate presentation where they lose accuracy. Also, full short channel models become too complex for hand calculations. Although the long channel models are also not accurate, they allow manual problem solving insight into the various bias regions of the transistors. We originally wrote the book using short channel models, but found the simplified analytical expressions clumsy and inaccurate. A second observation is that modern industry papers and oral presentations often refer to long channel models despite use of short channel transistors. The more accurate short channel models are best left to graduate courses and detailed computer models.

Other choices were made to avoid overly complex material at this undergraduate stage. For example, combinational logic power analysis uses the truth table analysis rather than logical effort. Chapter 9 on memory keeps the timing description simple, but to the point. Memory design deserves a whole book for adequate description.

The problems in this book most efficiently use the modern equation solving ability of scientific calculators. One great learning advantage is that time is spent on the problem itself and little on the grind of solving quadratic equations. An unknown variable can be embedded anywhere in the equation, and the scientific calculator doesn't care. It solves for the unknown variable in seconds. Students and instructors can solve these problems anyway they desire, but the scientific calculator is truly an advance of modern digital circuit products.

Author Background

This book reflects the experience of both authors who have taught this material at the graduate and undergraduate level, and who have stayed close to the digital electronic industry in their careers. Both authors did sabbaticals with the Intel Corp.; Prof. Segura at the Intel campus in Portland Oregon, and Prof. Hawkins in Rio Rancho, New Mexico. Prof. Segura also did sabbatical work at Philips Semiconductor as well as receiving numerous research contracts from industry. Prof. Hawkins worked closely with the Sandia National Laboratory in New Mexico for over 20 years in their CMOS integrated circuits group. Both authors have long histories of committee work for the European DATE conference, the International Test Conference, and the VLSI Test Symposium. Prof. Hawkins was editor of the Electron Device Failure Analysis Magazine. All of these activities outside of the classroom influenced our choice of material and style in the book. It is long overdue for electrical and computer engineering undergraduate students to rid themselves of bipolar logic circuits and receive a course dedicated to digital CMOS electronics.

A Suggested Semester Chapter Order

Chapter 1 Basic Logic Gate and Circuit Theory -1 week

Chapter 2 Semiconductor Physics -1 week

Chapter 3 MOSFET Transistors -2 weeks

Chapter 5 CMOS Inverter -1.5 weeks

Chapter 6 CMOS NAND, NOR, and Transmission Gates -1 week

Chapter 7 CMOS Design Styles -1.5 weeks

Chapter 11 CMOS Layout -1 week

Chapter 12 How Chips are Made -1 week

Chapter 4 Metal Interconnection Properties -1 week

Chapter 8 Sequential Logic Gate Design and Timing -2 weeks

Chapter 9 Memory Circuits -1 week

Chapter 10 FPGAs -0.5 week

Chapter 4 on metal interconnects logically fits with device descriptions, but it interrupts the flow of electronic circuitry, so it was put later. Chapters 11 and 12 continue the emphasis on circuitry and the IC before returning to Chapter 4.

Introduction

"Any sufficiently advanced technology is indistinguishable from magic"

- Arthur C. Clarke's Third Law

The goal of this book is to prepare you to contribute to the computer evolution in the 21st century. It is about the electronics that propels the incredible surge in human communication and knowledge capability. The foundation of a computer is the transistor. Computer electronics deals with transistor level behavior of circuits that perform all of the computer logic operations such as adding, multiplying, storing, comparing, and any operation described by Boolean equations. Billions of transistors and their wire connections are embedded in small, thin, rectangular silicon computer chips. The total wire connections on these tiny chips may be several miles in length, and power dissipation may range from a few microwatts to over 100 watts. The chip is also referred to as an integrated circuit (IC). Chips are complicated, and Electrical and Computer Engineers must understand computing at this circuit level.

Engineers face challenges. How would you blend digital circuit knowledge with computer architecture to design a chip? How fast do we want to clock the computer, and where do we start? How do you interface chips on a circuit board? How much heat from chip power loss can you stand –how do you minimize it? As a customer, how do you talk to a chip designer? When your first chips are returned from the factory to evaluate and something is wrong, where do you begin to solve the problem? Failures may be temperature or power supply dependent and not simple static Boolean errors. What skills and knowledge do you need to identify and correct these failures? Whether you are an engineer at the chip level, or you design at the higher board and system level, the solutions often reside with knowledge of chip properties at the transistor level.

A knowledge hierarchy exists in electronics. Semiconductor physics describes diode and transistor action using model equations that allow calculation of transistor circuit node voltages and path currents. Specific transistor configurations then form the different logic gates, such as the inverter, NAND, NOR, transmission gates, the D-flip-flop, and more complex combinational logics gates derived from arbitrary Boolean statements. These logic gates electronically perform the Boolean operations that define the computer, and we must understand their properties. What are the voltage, current, temperature, power drain, propagation delay time, and noise margins properties?

A master clock oscillator drives sequential circuits with pulses that synchronize data movement of the Boolean operations in the computer. The clock speed is an important parameter and often the first specification that a buyer looks at when shopping for a computer. The amount of computer memory may be the next question. Memory subcircuits are extensively built into the computer chips. What is a standard memory cell and how are memories organized? Modern computer chips may dedicate over 70% of the total transistor count to embedded memory. Memory embedded in the chip allows faster computing as opposed to sending signals back and forth to external memory chips mounted on a circuit board.

We might take our computer-based miracles for granted such as the Internet, cell phone magic, email, Google, automobile electronics, biomedical instrumentation, GPS, YouTube, instant news, weather, and sports, automobile electronics, and yes video games. You might ask, “Hasn’t it always been like this?” The answer is no –the applications didn’t really get rolling until the early 1990s, and all of these modern products depended on fast, cheap, and small computer chips.

Transistors and Computers – Until Death do They Part

To get a better sense of our subject, let us track electronic progress in digital computer development and then its role in the Internet. We see not only the march of computers to smaller, faster, and cheaper but also the fascinating interplay of diverse forces. The Internet did not grow in a vacuum and neither did computers.

The first computer circuit we are aware of was called the flip-flop by its English inventors, Eccles and Jordan, almost 100 years ago. A flip-flop remains stable in one of two voltage states until triggered to the other state by an external electrical pulse. The flip-flop stores a voltage state. Computers were not thought of at that time so the flip-flop remained dormant for many years. But today up to millions of flip-flops exist in every computer chip from the advanced Internet server chip to the chips in modern coffee makers or dishwashers. Flip-flops are at the heart of synchronizing data transfer.

In the late 1930s primitive computers combined Boolean algebra with mechanical switches to demonstrate simple computing machines. World War II sparked an interest in using computers for scientific calculations. The first vacuum tube computer was the ENIAC at the University of

Pennsylvania in 1946. By the standards of its day, the 100 kHz clock was fast. It weighed 30 tons, was 80 x 8.5 x 3.5 feet, and dissipated 150 kW of power. The old flip-flop was now an integral part of computer electronics. But the vacuum tube was a relatively large device requiring a glass enclosed vacuum and a heated metal filament. Tubes had poor reliability and were a challenge to cool. Something better was needed.

Bell Labs had a vision in the 1930s that a small, switching device could be constructed in a pure solid material. Bell Labs was thinking of replacing the slow, clunky mechanical relays in their telephone switching centers and not about computer development. In 1947, they struck gold with demonstration of a small, solid-state device called the transistor. About five years later, transistor computers emerged in production from several companies. Transistors were a giant step toward smaller, cooler, and more reliable computers. These computers used discrete (individual) transistors that were mounted in small metal cans, and were not the small, integrated circuit chips with billions of transistors that were to follow. These mainframe computers as they were called still required a cooled, dedicated room, but steady progress was made into the 1970s when another revolution occurred.

Actually several things happened at the transistor level. The first was a rapid transition from the original Bell Labs transistor called a bipolar junction transistor to a newer device called a MOSFET transistor. The MOSFET was blended in a unique design style called CMOS that was markedly cooler. The cooler CMOS allowed more transistors to be placed on a single chip without overheating thus increasing the computer functionality. CMOS also had the unusual property that if the transistor size were shrunk, the transistor would operate faster.

A third feature was that the smaller size of a CMOS transistor allowed more chips to be manufactured in a single operation than before, because the total chip size could now be reduced. More chips could be accommodated per process run, and that drives the cost down. Often industry left the chip the same size and just added more transistors to increase functionality.

A final feature is that if the small particle defects that kill the chips in a production run remain the same density, then packing more chips in the same area will increase the fraction of good chips, i.e., the yield. This gives a marked cost savings. CMOS has dominated computing chip design since about 1980, and CMOS technology today remains the focus of intense development.

It is a manufacturing miracle that next generation chips could be sold for a lower price if the next generation transistor was smaller. That was huge, and today you still pay about the same price (or less) for a personal computer as one that is a few years older. And these newer chips go faster and give more functionality while keeping the chip temperature under control. These CMOS features really fueled the development of computing chips. The reader should pause and dwell on the significance. What other product offers more dramatic performance each year for the same price or less.

Transistors and Computers –How Deep Can the Friendship Go

In the early 1970s, Intel brought out the first microprocessors, first the 4-bit and then the 8–bit. Product innovation leaped on these transistor level advances. In 1974, the MITS Corp. in Albuquerque, New Mexico USA offered the first personal computer, the PC. The MITS Altair 8800 was a primitive PC requiring code to be entered by toggle switches, but it had a video monitor, and a typewriter size. It had a 2 MHz clock and cost $498 assembled. It was also the first computer to be personally owned. It used a single microprorocessor chip, the Intel 8080, to perform the computer function and many engineers bought them out of curiosity. Interestingly Bill Gates of the new Microsoft Corp. in Albuquerque wrote BASIC for the MITS Altair PC. In 1977, Apple launched the Apple II PC for $1200. No one had ever had a computer at that price, size and capability, and especially one they could call their own. But the IBM PC launched in 1980 had more impact because it brought in the business sector. There was no looking back. Businesses were being freed from the tedium of the big, central computer room, and later travelers found they could do work on the road with the coming of laptops. The ubiquitous typewriter was on the way out.

The PC launched a revolution in information accessibility that could not have been imagined. Technology and novel business enterprises were beginning to move together. A partnering of technology, business enterprise, and government support at crucial points drove this revolution. But a monster enterprise called the Internet lay quietly awaiting its entrance.

In the 1960s and 1970s, Internet development was marching in the background to its own beat driven by engineers and scientists who wanted to use each other’s specialty computers across the country. It was government funding through its ARPA agency that allowed a mainframe computer from UCLA to use an interface unit, called the Interface Message Processor (IMP), to talk to a similar hookup at Stanford University in October of 1969. Long distance sharing of computer resources had happened. Messages, later called email, were exchanged but it was regarded then as a secondary feature and not a big deal. In fact the first Internet exchange was not widely publicized. The response was sort of, “Isn’t that nice that scientists and engineers can use each other’s computers, but that won’t affect my life.” What an understatement.

The next necessary development occurred in 1989 when PC manufacturers began bundling internal modems in the PC. The Internet was now open to anyone. Email grew at a tremendous rate as users found it a good business tool, and as true today it was just plain fun to use. The mouse and graphical displays were huge steps toward friendly computers. And computer chips doubled their speed and transistor density about every two years following what is called Moore's Law. Then SPAM, viruses, and hackers showed their ugly heads.   SPAM is expensive in system bandwidth and the required electrical energy generation to support its Internet hunger.

The Internet went global with introduction of the World Wide Web. "www" was a concept from CERN in Europe that was demonstrated in 1991. We now see "www" in our URL addresses. Browsers quickly followed with MOSAIC from the University of Illinois and the NETSCAPE browser from Netscape Corp. Yahoo and Microsoft entered the competition, and the famous browser wars were on. Two students from Stanford introduced Google in 1998 with a novel concept in searches that became so successful that Google is now a verb.

Although clearly visible in these early applications, it was the special talents of the business entrepreneurs that carried the World Wide Web into its most recent surge. The incredible innovations now seem endless. The list of Amazon, eBay, PayPal, Google, Wikipedia, YouTube, bloggers, Refdesk, email delivery each morning of your favorite newspaper, instant check on stocks and weather, instant Google satellite maps of the Earth, on-line business carried on across the globe, twitter, and many more brings us to our present state of information availability. These applications required computer chips that were faster, smaller, and cheaper. The miracle applications needed the base technologies.

Computers –Is There a Limit

Computer chips depend on many disciplines. Electrical and computer engineers, computer scientists, mathematicians, physicists, chemical engineers, chemists, mechanical engineers, statisticians, manufacturing engineers, and marketing people work in harmony to achieve these miracle products. Technology products typically develop from an idea and a prototype demonstration. If the idea is sound, the product undergoes continuous improvement until performance limits are reached. How far can we push performance? Let's look at three other technologies to see their performance trajectory and how that might provide clues to our electronic future.

Train development spanned an increasing performance era from about 1820 to the 1950s. Then, except for the bullet train, it was over. Automobile development spanned from the 1890s to about the 1960s. Speed, comfort, and engine power peaked for trains and automobiles. Commercial aircraft basically spanned from 1903 to the 1960s when the Boeing 747 was produced. Speed and passenger carrying capability maxed out. Later in all three areas, the integrated circuit caused a second revolution in the 1980s. If you lived in the height of these technology rushes, there was a feeling that "progress" would never end. But the basic speed, power, and transport capability did end. What does this imply for computers?

Will the CMOS computer technology rush end? Will our computers provide dramatically more functionality each year? For several reasons, we believe that CMOS technology will hit a performance limit. If history is a guide, we will then squeeze every last design and manufacturing detail from our chips, and improvements will lie in more efficient manufacturing and using multiple processors on one chip. When will we see that soft end point? We see signs of reaching some of the limits now, so we hazard an educated guess that the CMOS performance limit may be reached by 2020. That is a guess. The exact date is immaterial to the thought that a performance limit exists within our professional lifetimes.

When CMOS performance development ends, we expect research will continue seeking another manufacturable electronic technology. The urge is strong to build faster, smaller, cheaper computers, and it will be novel transistor or transistor like devices that pushes us even further.

One significant challenge deals with electrical power. Tom Friedman in his book "Hot, Flat, and Crowded" (p. 31) quoted a Sun Microsystem engineer who put future power demand in perspective. He observed that the earth will add about one billion persons in the next 12 years. If each person were given a 60-watt light bulb then 60 billion watts of power generation would be required. If each person used that bulb for an average of four hours per day, then that average 10 billion watts would require about 20 coal-fired power plants. If each of the billion persons were also allowed to use a 120-watt computer for four hours then we would need an additional 40 power plants. The earth is power limited, and the pressure on low power computer chip design is huge.

The concept of technology is little different than early mankind using a wheel to support a cart to carry heavy weights. Technology is the use of materials and natural laws to ease our burdens. Electronic technology is no different, but we know that all technology has a downside. The Internet has done miracles, but that hasn't stopped hackers, Spammers, and swindlers from peddling their dark objectives. The Internet can bring instant and accurate news, but it can also bring instant and inaccurate propaganda. These are issues to deal with as it has always been with technology. We should keep our eye on the benefits and continue the historic human battle of fighting the misuse of technology.

Future

We won't speculate much on the future other than that there is one. The Internet brought a speed to changes in business and technology that we now take for granted. Startling new products will appear, and some old product names will disappear. This book addresses the education of the next generation of engineers who will continue to move this historic epic in information accessibility.

Acknowledgements: The Internet itself is a rapid source of fine details of computer history. Also, an excellent article appeared in Vanity Fair Magazine in August 2008, "How the Web was Won: An Oral History of the Internet" page 96, by K. Mayo and P. Newcomb. Thomas Friedman's, "Hot Flat, and Crowded" is an excellent book of what the future may have in store for energy and technology (Farrar, Straus, and Giroux Publisher, 2008).