

THE ESSENTIALS OF **Computer Organization and Architecture**

FIFTH EDITION

Linda Null, PhD

The Pennsylvania State University

Julia Lobur



JONES & BARTLETT
LEARNING



World Headquarters
Jones & Bartlett Learning
5 Wall Street
Burlington, MA 01803
978-443-5000
info@jblearning.com
www.jblearning.com

Jones & Bartlett Learning books and products are available through most bookstores and online booksellers. To contact Jones & Bartlett Learning directly, call 800-832-0034, fax 978-443-8000, or visit our website, www.jblearning.com.

Substantial discounts on bulk quantities of Jones & Bartlett Learning publications are available to corporations, professional associations, and other qualified organizations. For details and specific discount information, contact the special sales department at Jones & Bartlett Learning via the above contact information or send an email to specialsales@jblearning.com.

Copyright © 2019 by Jones & Bartlett Learning, LLC, an Ascend Learning Company

All rights reserved. No part of the material protected by this copyright may be reproduced or utilized in any form, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the copyright owner.

The content, statements, views, and opinions herein are the sole expression of the respective authors and not that of Jones & Bartlett Learning, LLC. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement or recommendation by Jones & Bartlett Learning, LLC and such reference shall not be used for advertising or product endorsement purposes. All trademarks displayed are the trademarks of the parties noted herein. *Essentials of Computer Organization and Architecture, Fifth Edition* is an independent publication and has not been authorized, sponsored, or otherwise approved by the owners of the trademarks or service marks referenced in this product.

There may be images in this book that feature models; these models do not necessarily endorse, represent, or participate in the activities represented in the images. Any screenshots in this product are for educational and instructive purposes only. Any individuals and scenarios featured in the case studies throughout this product may be real or fictitious, but are used for instructional purposes only.

13685-2

Production Credits

VP, Product Management: David D. Cella
Director of Product Management: Matthew Kane
Product Manager: Laura Pagluica
Product Assistant: Rebecca Feeney
Associate Production Editor: Alex Schab
Director of Marketing: Andrea DeFronzo
Production Services Manager: Colleen Lamy
Product Fulfillment Manager: Wendy Kilborn

Composition: codeMantra U.S. LLC
Cover Design: Kristin E. Parker
Rights & Media Specialist: Thais Miller
Media Development Editor: Shannon Sheehan
Cover Image (Title Page, Part Opener, Chapter Opener): © Adam Petto/Getty Images and © antos777/Getty Images
Printing and Binding: Edwards Brothers Malloy
Cover Printing: Edwards Brothers Malloy

Library of Congress Cataloging-in-Publication Data

Names: Null, Linda, author.
Title: *Essentials of computer organization and architecture* / Linda Null.
Description: Fifth edition. | Burlington, Massachusetts: Jones & Bartlett Learning, 2018. | Includes bibliographical references and index.
Identifiers: LCCN 2017044949 | ISBN 9781284123036 (hardback)
Subjects: LCSH: Computer organization. | Computer architecture. |
BISAC: COMPUTERS / Software Development & Engineering / General.
Classification: LCC QA76.9.C643 N85 2018 | DDC 004.2/2—dc23
LC record available at <https://lccn.loc.gov/2017044949>

6048

Printed in the United States of America
22 21 20 19 18 10 9 8 7 6 5 4 3 2 1

In memory of my husband, Tim Wahls, who exemplified the beauty and wonder in everything around me and taught me that love does not end with death. How amazingly lucky I was to have someone who made saying goodbye so difficult,

—L. M. N.

To the loving memory of my mother, Anna J. Surowski, who made all things possible for her girls.

—J. M. L.



Contents

Preface

xv

CHAPTER

1

Introduction

1

1.1 Overview 1

1.2 Computer Systems 3

1.2.1 The Main Components of a Computer 3

1.2.2 System Components 4

1.2.3 Classification of Computing Devices 6

1.3 An Example System: Wading Through the Jargon 6

1.4 Standards Organizations 22

1.5 Historical Development 23

1.5.1 Generation Zero: Mechanical Calculating Machines (1642–1945) 24

1.5.2 The First Generation: Vacuum Tube Computers (1945–1953) 26

1.5.3 The Second Generation: Transistorized Computers (1954–1965) 30

1.5.4 The Third Generation: Integrated Circuit Computers (1965–1980) 31

1.5.5 The Fourth Generation: VLSI Computers (1980–????) 32

1.5.6 Moore's Law 35

1.6 The Computer Level Hierarchy 36

1.7 Cloud Computing: Computing as a Service 39

1.8 The Fragility of the Internet 43

1.9 The von Neumann Model 45

1.10 Non–von Neumann Models 48

1.11 Parallel Processors and Parallel Computing 49

Chapter Summary 54

Further Reading 54

References 55

v

Review of Essential Terms and Concepts	56
Exercises	58

CHAPTER 2

Data Representation in Computer Systems

61

2.1	Introduction	61
2.2	Positional Numbering Systems	62
2.3	Converting Between Bases	62
2.3.1	Converting Unsigned Whole Numbers	63
2.3.2	Converting Fractions	65
2.3.3	Converting Between Power-of-Two Radices	68
2.4	Signed Integer Representation	68
2.4.1	Signed Magnitude	68
2.4.2	Complement Systems	74
2.4.3	Excess-M Representation for Signed Numbers	81
2.4.4	Unsigned Versus Signed Numbers	82
2.4.5	Computers, Arithmetic, and Booth's Algorithm	83
2.4.6	Carry Versus Overflow	86
2.4.7	Binary Multiplication and Division Using Shifting	87
2.5	Floating-Point Representation	89
2.5.1	A Simple Model	90
2.5.2	Floating-Point Arithmetic	93
2.5.3	Floating-Point Errors	94
2.5.4	The IEEE-754 Floating-Point Standard	95
2.5.5	Range, Precision, and Accuracy	98
2.5.6	Additional Problems with Floating-Point Numbers	99
2.6	Character Codes	102
2.6.1	Binary-Coded Decimal	102
2.6.2	EBCDIC	104
2.6.3	ASCII	105
2.6.4	Unicode	105
2.7	Error Detection and Correction	109
2.7.1	Cyclic Redundancy Check	109
2.7.2	Hamming Codes	112
2.7.3	Reed-Solomon	119
	Chapter Summary	120
	Further Reading	120
	References	121

Review of Essential Terms and Concepts	121
Exercises	123

CHAPTER 3

Boolean Algebra and Digital Logic

135

3.1	Introduction	135
3.2	Boolean Algebra	136
3.2.1	Boolean Expressions	137
3.2.2	Boolean Identities	138
3.2.3	Simplification of Boolean Expressions	140
3.2.4	Complements	143
3.2.5	Representing Boolean Functions	144
3.3	Logic Gates	146
3.3.1	Symbols for Logic Gates	146
3.3.2	Universal Gates	147
3.3.3	Multiple Input Gates	148
3.4	Karnaugh Maps	149
3.4.1	Introduction	149
3.4.2	Description of Kmaps and Terminology	149
3.4.3	Kmap Simplification for Two Variables	151
3.4.4	Kmap Simplification for Three Variables	153
3.4.5	Kmap Simplification for Four Variables	156
3.4.6	Don't Care Conditions	159
3.4.7	Summary	160
3.5	Digital Components	160
3.5.1	Digital Circuits and Their Relationship to Boolean Algebra	160
3.5.2	Integrated Circuits	162
3.5.3	Putting It All Together: From Problem Description to Circuit	164
3.6	Combinational Circuits	166
3.6.1	Basic Concepts	166
3.6.2	Examples of Typical Combinational Circuits	166
3.7	Sequential Circuits	175
3.7.1	Basic Concepts	176
3.7.2	Clocks	176
3.7.3	Flip-Flops	176
3.7.4	Finite-State Machines	180
3.7.5	Examples of Sequential Circuits	188
3.7.6	An Application of Sequential Logic: Convolutional Coding and Viterbi Detection	194

3.8	Designing Circuits	199
	Chapter Summary	201
	Further Reading	202
	References	204
	Review of Essential Terms and Concepts	204
	Exercises	205

CHAPTER	MARIE: An Introduction to a Simple Computer	219
4		
4.1	Introduction	219
4.2	CPU Basics and Organization	219
	4.2.1 The Registers	220
	4.2.2 The ALU	221
	4.2.3 The Control Unit	221
4.3	The Bus	221
4.4	Clocks	225
4.5	The Input/Output Subsystem	227
4.6	Memory Organization and Addressing	227
4.7	Interrupts	235
4.8	MARIE	236
	4.8.1 The Architecture	236
	4.8.2 Registers and Buses	236
	4.8.3 Instruction Set Architecture	238
	4.8.4 Register Transfer Notation	242
4.9	Instruction Processing	244
	4.9.1 The Fetch–Decode–Execute Cycle	244
	4.9.2 Interrupts and the Instruction Cycle	246
	4.9.3 MARIE’s I/O	249
4.10	A Simple Program	249
4.11	A Discussion on Assemblers	252
	4.11.1 What Do Assemblers Do?	252
	4.11.2 Why Use Assembly Language?	254
4.12	Extending Our Instruction Set	255
4.13	A Discussion on Decoding: Hardwired Versus Microprogrammed Control	262
	4.13.1 Machine Control	262
	4.13.2 Hardwired Control	265
	4.13.3 Microprogrammed Control	270

4.14	Real-World Examples of Computer Architectures	274
4.14.1	Intel Architectures	275
4.14.2	MIPS Architectures	282
	Chapter Summary	284
	Further Reading	286
	References	286
	Review of Essential Terms and Concepts	288
	Exercises	289
	True or False	298

CHAPTER 5

A Closer Look at Instruction Set Architectures 299

5.1	Introduction	299
5.2	Instruction Formats	299
5.2.1	Design Decisions for Instruction Sets	300
5.2.2	Little Versus Big Endian	301
5.2.3	Internal Storage in the CPU: Stacks Versus Registers	304
5.2.4	Number of Operands and Instruction Length	305
5.2.5	Expanding Opcodes	311
5.3	Instruction Types	316
5.3.1	Data Movement	316
5.3.2	Arithmetic Operations	316
5.3.3	Boolean Logic Instructions	316
5.3.4	Bit Manipulation Instructions	317
5.3.5	Input/Output Instructions	317
5.3.6	Instructions for Transfer of Control	318
5.3.7	Special-Purpose Instructions	318
5.3.8	Instruction Set Orthogonality	318
5.4	Addressing	319
5.4.1	Data Types	319
5.4.2	Address Modes	319
5.5	Instruction Pipelining	322
5.6	Real-World Examples of ISAs	328
5.6.1	Intel	328
5.6.2	MIPS	329
5.6.3	Java Virtual Machine	333
5.6.4	ARM	337
	Chapter Summary	340

Further Reading	341
References	342
Review of Essential Terms and Concepts	343
Exercises	344
True or False	349

CHAPTER 6

Memory

351

6.1	Introduction	351
6.2	Types of Memory	351
6.3	The Memory Hierarchy	353
6.3.1	Locality of Reference	356
6.4	Cache Memory	356
6.4.1	Cache Mapping Schemes	359
6.4.2	Replacement Policies	374
6.4.3	Effective Access Time and Hit Ratio	377
6.4.4	When Does Caching Break Down?	378
6.4.5	Cache Write Policies	379
6.4.6	Instruction and Data Caches	379
6.4.7	Levels of Cache	382
6.5	Virtual Memory	383
6.5.1	Paging	384
6.5.2	Effective Access Time Using Paging	392
6.5.3	Putting It All Together: Using Cache, TLBs, and Paging	393
6.5.4	Advantages and Disadvantages of Paging and Virtual Memory	395
6.5.5	Segmentation	396
6.5.6	Paging Combined with Segmentation	397
6.6	Real-World Examples of Memory Management	397
	Chapter Summary	399
	Further Reading	400
	References	401
	Review of Essential Terms and Concepts	401
	Exercises	402

**CHAPTER
7****Input/Output Systems****411**

- 7.1 Introduction 411**
- 7.2 I/O and Performance 412**
- 7.3 Amdahl's Law 412**
- 7.4 I/O Architectures 415**
 - 7.4.1 I/O Control Methods 417
 - 7.4.2 Character I/O Versus Block I/O 425
 - 7.4.3 I/O Bus Operation 425
 - 7.4.4 I/O Buses and Interfaces 429
- 7.5 Data Transmission Modes 430**
 - 7.5.1 Parallel Data Transmission 430
 - 7.5.2 Serial Data Transmission 433
- 7.6 Disk Technology 434**
 - 7.6.1 Rigid Disk Drives 435
 - 7.6.2 Solid State Drives 439
- 7.7 Optical Disks 442**
 - 7.7.1 CD-ROM 442
 - 7.7.2 DVD 446
 - 7.7.3 Blue-Violet Laser Discs 447
 - 7.7.4 Optical Disk Recording Methods 448
- 7.8 Magnetic Tape 449**
 - 7.8.1 LTO: Linear Tape-Open 450
- 7.9 RAID 453**
 - 7.9.1 RAID Level 0 453
 - 7.9.2 RAID Level 1 454
 - 7.9.3 RAID Level 2 455
 - 7.9.4 RAID Level 3 456
 - 7.9.5 RAID Level 4 457
 - 7.9.6 RAID Level 5 458
 - 7.9.7 RAID Level 6 458
 - 7.9.8 RAID DP 460
 - 7.9.9 Hybrid RAID Systems 462
- 7.10 The Future of Data Storage 464**
- Chapter Summary 468
- Further Reading 469
- References 470
- Review of Essential Terms and Concepts 471
- Exercises 472

CHAPTER	System Software	481
8		
8.1	Introduction 481	
8.2	Operating Systems 482	
8.2.1	Operating Systems History 483	
8.2.2	Operating System Design 488	
8.2.3	Operating System Services 490	
8.3	Protected Environments 494	
8.3.1	Virtual Machines 495	
8.3.2	Subsystems and Partitions 498	
8.3.3	Protected Environments and the Evolution of Systems Architectures 500	
8.4	Programming Tools 501	
8.4.1	Assemblers and Assembly 502	
8.4.2	Link Editors 505	
8.4.3	Dynamic Link Libraries 506	
8.4.4	Compilers 508	
8.4.5	Interpreters 512	
8.5	Java: All of the Above 513	
	Chapter Summary 521	
	Further Reading 522	
	References 523	
	Review of Essential Terms and Concepts 524	
	Exercises 525	
CHAPTER	Alternative Architectures	527
9		
9.1	Introduction 527	
9.2	RISC Machines 528	
9.3	Flynn's Taxonomy 536	
9.4	Parallel and Multiprocessor Architectures 539	
9.4.1	Superscalar and VLIW 540	
9.4.2	Vector Processors 542	
9.4.3	Interconnection Networks 543	
9.4.4	Shared Memory Multiprocessors 548	
9.4.5	Distributed Computing 552	
9.5	Alternative Parallel Processing Approaches 556	
9.5.1	Dataflow Computing 556	
9.5.2	Neural Networks 559	
9.5.3	Systolic Arrays 562	

9.6 Quantum Computing 563

Chapter Summary 567

Further Reading 568

References 569

Review of Essential Terms and Concepts 571

Exercises 572

**CHAPTER
10**

Topics in Embedded Systems

577

10.1 Introduction 577

10.2 An Overview of Embedded Hardware 579

10.2.1 Off-the-Shelf Embedded System Hardware 579

10.2.2 Configurable Hardware 583

10.2.3 Custom-Designed Embedded Hardware 590

10.3 An Overview of Embedded Software 599

10.3.1 Embedded Systems Memory Organization 599

10.3.2 Embedded Operating Systems 600

10.3.3 Embedded Systems Software Development 604

Chapter Summary 606

Further Reading 608

References 609

Review of Essential Terms and Concepts 610

Exercises 611

**CHAPTER
11**

Performance Measurement and Analysis

613

11.1 Introduction 613

11.2 Computer Performance Equations 614

11.3 Mathematical Preliminaries 615

11.3.1 What the Means Mean 616

11.3.2 The Statistics and Semantics 622

11.4 Benchmarking 623

11.4.1 Clock Rate, MIPS, and FLOPS 625

11.4.2 Synthetic Benchmarks: Whetstone, Linpack, and Dhrystone 626

11.4.3 Standard Performance Evaluation Corporation Benchmarks 628

11.4.4 Transaction Processing Performance Council Benchmarks 630

11.4.5 System Simulation 634

11.5	CPU Performance Optimization	637
11.5.1	Branch Optimization	637
11.5.2	Use of Good Algorithms and Simple Code	640
11.6	Disk Performance	643
11.6.1	Understanding the Problem	644
11.6.2	Physical Considerations	645
11.6.3	Logical Considerations	646
	Chapter Summary	651
	Further Reading	652
	References	653
	Review of Essential Terms and Concepts	654
	Exercises	655

Glossary	659
-----------------	------------

Answers and Hints for Selected Exercises	701
---	------------

Index	715
--------------	------------

Preface

TO THE STUDENT

This is a book about computer organization and architecture. It focuses on the function and design of the various components necessary to process information digitally. We present computing systems as a series of layers, starting with low-level hardware and progressing to higher-level software, including assemblers and operating systems. These levels constitute a hierarchy of virtual machines. The study of computer organization focuses on this hierarchy and the issues involved with how we partition the levels and how each level is implemented. The study of computer architecture focuses on the interface between hardware and software, and emphasizes the structure and behavior of the system. The majority of information contained in this textbook is devoted to computer hardware, computer organization and architecture, and their relationship to software performance.

Students invariably ask, “Why, if I am a computer science major, must I learn about computer hardware? Isn’t that for computer engineers? Why do I care what the inside of a computer looks like?” As computer users, we probably do not have to worry about this any more than we need to know what our cars look like under the hood in order to drive them. We can certainly write high-level language programs without understanding how these programs execute; we can use various application packages without understanding how they really work. But what happens when the program we have written needs to be faster and more efficient, or the application we are using doesn’t do precisely what we want? As computer

scientists, we need a basic understanding of the computer system itself in order to rectify these problems.

There is a fundamental relationship between the computer hardware and the many aspects of programming and software components in computer systems. In order to write good software, it is very important to understand the computer system as a whole. Understanding hardware can help you explain the mysterious errors that sometimes creep into your programs, such as the infamous segmentation fault or bus error. The level of knowledge about computer organization and computer architecture that a high-level programmer must have depends on the task the high-level programmer is attempting to complete.

For example, to write compilers, you must understand the particular hardware to which you are compiling. Some of the ideas used in hardware (such as pipelining) can be adapted to compilation techniques, thus making the compiler faster and more efficient. To model large, complex, real-world systems, you must understand how floating-point arithmetic should work, and how it does work (which are not necessarily the same thing). To write device drivers for video, disks, or other I/O devices, you need a good understanding of I/O interfacing and computer architecture in general. If you want to work on embedded systems, which are usually very resource constrained, you must understand all of the time, space, and price trade-offs. To do research on, and make recommendations for, hardware systems, networks, or specific algorithms, you must acquire an understanding of benchmarking and then learn how to present performance results adequately. Before buying hardware, you need to understand benchmarking and all the ways that others can *manipulate* the performance results to “prove” that one system is better than another. Regardless of our particular area of expertise, as computer scientists, it is imperative that we understand how hardware interacts with software.

You may also be wondering why a book with the word *essentials* in its title is so large. The reason is twofold. First, the subject of computer organization is expansive and it grows by the day. Second, there is little agreement as to which topics from within this burgeoning sea of information are truly essential and which are just helpful to know. In writing this book, one goal was to provide a concise text compliant with the computer architecture curriculum guidelines jointly published by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronic Engineers (IEEE). These guidelines encompass the subject matter that experts agree constitutes the “essential” core body of knowledge relevant to the subject of computer organization and architecture.

We have augmented the ACM/IEEE recommendations with subject matter that we feel is useful—if not essential—to your continuing computer science studies and to your professional advancement. The topics that we feel will help you in your continuing computer science studies include operating systems, compilers, database management, and data communications. Other subjects are included because they will help you understand how actual systems work in real life.

We hope that you find reading this book an enjoyable experience, and that you take time to delve deeper into some of the material we have presented. It is our intention that this book will serve as a useful reference long after your formal course is complete. Although we give you a substantial amount of information, it is only a foundation upon which you can build throughout the remainder of your studies and your career. Successful computer professionals continually add to their knowledge about how computers work. Welcome to the start of your journey.

TO THE INSTRUCTOR

This book is the outgrowth of two computer science organization and architecture classes taught at Penn State Harrisburg. As the computer science curriculum evolved, we found it necessary not only to modify the material taught in the courses, but also to condense the courses from a two-semester sequence into a three-credit, one-semester course. Many other schools have also recognized the need to compress material in order to make room for emerging topics. This new course, as well as this textbook, is primarily for computer science majors and is intended to address the topics in computer organization and architecture with which computer science majors must be familiar. This book not only integrates the underlying principles in these areas, but it also introduces and motivates the topics, providing the breadth necessary for majors while providing the depth necessary for continuing studies in computer science.

Our primary objective in writing this book was to change the way computer organization and architecture are typically taught. A computer science major should leave a computer organization and architecture class with not only an understanding of the important general concepts on which the digital computer is founded, but also with a comprehension of how those concepts apply to the real world. These concepts should transcend vendor-specific terminology and design; in fact, students should be able to take concepts given in the specific and translate to the generic and vice versa. In addition, students must develop a firm foundation for further study in the major.

The title of our book, *The Essentials of Computer Organization and Architecture*, is intended to convey that the topics presented in the text are those for which every computer science major should have exposure, familiarity, or mastery. We do not expect students using our textbook to have complete mastery of all topics presented. It is our firm belief, however, that there are certain topics that must be mastered; there are those topics about which students must have a definite familiarity; and there are certain topics for which a brief introduction and exposure are adequate.

We do not feel that concepts presented in sufficient depth can be learned by studying general principles in isolation. We therefore present the topics as an integrated set of solutions, not simply a collection of individual pieces of information. We feel our explanations, examples, exercises, tutorials, and simulators all

combine to provide the student with a total learning experience that exposes the inner workings of a modern digital computer at the appropriate level.

We have written this textbook in an informal style, omitting unnecessary jargon, writing clearly and concisely, and avoiding unnecessary abstraction, in hopes of increasing student enthusiasm. We have also broadened the range of topics typically found in a first-level architecture book to include system software, a brief tour of operating systems, performance issues, alternative architectures, and a concise introduction to networking, as these topics are intimately related to computer hardware. Like most books, we have chosen an architectural model, but it is one that we have designed with simplicity in mind.

Relationship to CS2013

In October 2013, the ACM/IEEE Joint Task Force unveiled Computer Science Curricula 2013 (CS2013). Although we are primarily concerned with the Computer Architecture knowledge area, these guidelines suggest integrating the core knowledge throughout the curriculum. Therefore, we also call attention to additional knowledge areas beyond architecture that are addressed in this book.

CS2013 is a comprehensive revision of CS2008, mostly the result of focusing on the *essential concepts* in the computer science curriculum while still being flexible enough to meet individual institutional needs. These guidelines introduce the notion of Core Tier-1 and Core Tier-2 topics, in addition to elective topics. Core Tier-1 topics are those that should be part of every computer science curriculum. Core Tier-2 topics are those that are considered essential enough that a computer science curriculum should contain 90–100% of these topics. Elective topics are those that allow curricula to provide breadth and depth. The suggested coverage for each topic is listed in lecture hours.

The main change in the Architecture and Organization (AR) knowledge area from CS2008 to CS2013 is a reduction of lecture hours from 36 to 16; however, a new area, System Fundamentals (SF), has been introduced and includes some concepts previously found in the AR module (including hardware building blocks and architectural organization). The interested reader is referred to the CS2013 guidelines (<http://www.acm.org/education/curricula-recommendations>) for more information on what the individual knowledge areas include.

We are pleased that the fifth edition of *The Essentials of Computer Organization and Architecture* is in direct correlation with the ACM/IEEE CS2013 guidelines for computer organization and architecture, in addition to integrating material from additional knowledge units. Table P.1 indicates which chapters of this textbook satisfy the eight topics listed in the AR knowledge area. For the other knowledge areas, only the topics that are covered in this textbook are listed.

AR – Architecture	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
Digital Logic and Digital Systems		3	N	1, 3, 4
Machine-Level Representation of Data		3	N	1, 2
Assembly-Level Machine Organization		6	N	1, 4, 5, 7, 8, 9
Memory System Organization and Arch		3	N	2, 6, 7, 13
Interfacing and Communication		1	N	4, 7, 12
Functional Organization			Y	4, 5
Multiprocessing and Alternative Archs			Y	9
Performance Enhancements			Y	9, 11
NC – Networking and Communication	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
Introduction	1.5		N	12
Networked Applications	1.5		N	12
Reliable Data Delivery		2	N	12
Routing and Forwarding		1.5	N	12
OS – Operating Systems	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
Overview of Operating Systems	2		N	8
Memory Management		3	N	6
Virtual Machines			Y	8
File Systems			Y	7
Real-Time and Embedded Systems			Y	10
System Performance Evaluations			Y	6, 11
PD – Parallel and Distributed Computing	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
Parallel Architecture	1	1	N	9
Distributed Systems			Y	9
Cloud Computing			Y	1, 9, 13
SF – Systems Fundamentals	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
Computational Paradigms	3		N	3, 4, 9
State and State Machines	6		N	3
Parallelism	3		N	9
Evaluation	3		N	11
Proximity		3	N	6
SP – Social Issues and Professional Practice	Core Tier 1 Hours	Core Tier 2 Hours	Includes Electives	Chapters
History			Y	1

TABLE P.1 ACM/IEEE CS2013 Topics Covered in This Book

Why Another Text?

No one can deny there is a plethora of textbooks for teaching computer organization and architecture already on the market. In our 40-plus years of teaching these courses, we have used many very good textbooks. However, each time we have taught the course, the content has evolved, and eventually, we discovered we were writing significantly more course notes to bridge the gap between the material in the textbook and the material we deemed necessary to present in our classes. We found that our course material was migrating from a computer engineering approach to organization and architecture toward a computer science approach to these topics. When the decision was made to fold the organization class and the architecture class into one course, we simply could not find a textbook that covered the material we felt was necessary for our majors, written from a computer science point of view, written without machine-specific terminology, and designed to motivate the topics before covering them.

In this textbook, we hope to convey the spirit of design used in the development of modern computing systems and what effect this has on computer science students. Students, however, must have a strong understanding of the basic concepts before they can understand and appreciate the intangible aspects of design. Most organization and architecture textbooks present a similar subset of technical information regarding these basics. We, however, pay particular attention to the level at which the information should be covered, and to presenting that information in the context that has relevance for computer science students. For example, throughout this book, when concrete examples are necessary, we offer examples for personal computers, enterprise systems, and mainframes, as these are the types of systems most likely to be encountered. We avoid the PC bias prevalent in similar books in the hope that students will gain an appreciation for the differences and similarities between various platforms, as well as the roles they play in today's automated infrastructures. Too often, textbooks forget that motivation is, perhaps, the single most important key in learning. To that end, we include many real-world examples, while attempting to maintain a balance between theory and application.

Features

We have included many features in this textbook to emphasize the various concepts in computer organization and architecture, and to make the material more accessible to students. Some of the features are:

- *Sidebars.* These sidebars include interesting tidbits of information that go a step beyond the main focus of the chapter, thus allowing readers to delve further into the material. New to this edition, we include “Hints and Tips” sidebars focusing on potential problem areas.
- *Real-World Examples.* We have integrated the textbook with examples from real life to give students a better understanding of how technology and techniques are combined for practical purposes.
- *Chapter Summaries.* These sections provide brief yet concise summaries of the main points in each chapter.

- *Further Reading.* These sections list additional sources for those readers who wish to investigate any of the topics in more detail. They contain references to definitive papers and books related to the chapter topics.
- *Review Questions.* Each chapter contains a set of review questions designed to ensure that the reader has a firm grasp of the material.
- *Chapter Exercises.* Each chapter has a broad selection of exercises to reinforce the ideas presented. More challenging exercises are marked with an asterisk.
- *Answers to Selected Exercises.* To ensure that students are on the right track, we provide answers to representative questions from each chapter. Questions with answers in the back of the text are marked with a blue diamond.
- *Appendix.* The appendix provides a brief introduction or review of data structures, including topics such as stacks, linked lists, and trees.
- *Glossary.* An extensive glossary includes brief definitions of all key terms from the chapters.
- *Index.* An exhaustive index is provided with this book, with multiple cross-references, to make finding terms and concepts easier for the reader.

About the Authors

We bring to this textbook not only 40-plus years of combined teaching experience, but also 30-plus years of industry experience. Our combined efforts therefore stress the underlying principles of computer organization and architecture and how these topics relate in practice. We include real-life examples to help students appreciate how these fundamental concepts are applied in the world of computing.

Linda Null holds a PhD in computer science from Iowa State University, an MS in computer science from Iowa State University, an MS in computer science education from Northwest Missouri State University, an MS in mathematics education from Northwest Missouri State University, and a BS in mathematics and English from Northwest Missouri State University. She has been teaching mathematics and computer science for more than 40 years and is currently the computer science graduate program coordinator and associate program chair at the Pennsylvania State University Harrisburg campus, where she has been a member of the faculty since 1995. She has received numerous teaching awards, including the Penn State University Teaching Fellow Award, the Penn State Harrisburg Teaching Excellence Award, and the Penn State Harrisburg School of Science, Engineering, and Technology Excellence in Teaching Award. Her areas of interest include computer organization and architecture, operating systems, computer science education, and computer security.

Julia Lobur has been a practitioner in the computer industry for more than 30 years. She has held positions as systems consultant, staff programmer/analyst, systems and network designer, software development manager, and project

manager, in addition to part-time teaching duties. Julia holds an MS in computer science and is an IEEE Certified Software Development Professional.

Prerequisites

The typical background necessary for a student using this textbook includes a year of programming experience using a high-level procedural language. Students are also expected to have taken a year of college-level mathematics (calculus or discrete mathematics), as this textbook assumes and incorporates these mathematical concepts. This book assumes no prior knowledge of computer hardware.

A computer organization and architecture class is customarily a prerequisite for an undergraduate operating systems class (students must know about the memory hierarchy, concurrency, exceptions, and interrupts), compilers (students must know about instruction sets, memory addressing, and linking), networking (students must understand the hardware of a system before attempting to understand the network that ties these components together), and of course, any advanced architecture class. This text covers the topics necessary for these courses.

General Organization and Coverage

Our presentation of concepts in this textbook is an attempt at a concise yet thorough coverage of the topics we feel are essential for the computer science major. We do not feel the best way to do this is by “compartmentalizing” the various topics; therefore, we have chosen a structured yet integrated approach where each topic is covered in the context of the entire computer system.

As with many popular texts, we have taken a bottom-up approach, starting with the digital logic level and building to the application level, which students should be familiar with before starting the class. The text is carefully structured so that the reader understands one level before moving on to the next. By the time the reader reaches the application level, all the necessary concepts in computer organization and architecture have been presented. Our goal is to allow the students to tie the hardware knowledge covered in this book to the concepts learned in their introductory programming classes, resulting in a complete and thorough picture of how hardware and software fit together. Ultimately, the extent of hardware understanding has a significant influence on software design and performance. If students can build a firm foundation in hardware fundamentals, this will go a long way toward helping them to become better computer scientists.

The concepts in computer organization and architecture are integral to many of the everyday tasks that computer professionals perform. To address the numerous areas in which a computer professional should be educated, we have taken a high-level look at computer architecture, providing low-level coverage only when deemed necessary for an understanding of a specific concept. For example, when discussing ISAs, many hardware-dependent issues are introduced in the context of different case studies to both differentiate and reinforce the issues associated with ISA design.

The text is divided into 13 chapters and an appendix, as follows:

- **Chapter 1** provides a historical overview of computing in general, pointing out the many milestones in the development of computing systems and allowing the reader to visualize how we arrived at the current state of computing. This chapter introduces the necessary terminology, the basic components in a computer system, the various logical levels of a computer system, and the von Neumann computer model. It provides a high-level view of the computer system, as well as the motivation and necessary concepts for further study.
- **Chapter 2** provides thorough coverage of the various means computers use to represent both numerical and character information. Addition, subtraction, multiplication, and division are covered once the reader has been exposed to number bases and the typical numeric representation techniques, including one's complement, two's complement, and BCD. In addition, EBCDIC, ASCII, and Unicode character representations are addressed. Fixed- and floating-point representation are also introduced. Codes for data recording and error detection and correction are covered briefly. Codes for data transmission and recording are described in a special "Focus On" section.
- **Chapter 3** is a classic presentation of digital logic and how it relates to Boolean algebra. This chapter covers both combinational and sequential logic in sufficient detail to allow the reader to understand the logical makeup of more complicated MSI (medium-scale integration) circuits (such as decoders). More complex circuits, such as buses and memory, are also included. We have included optimization and Kmaps in a special "Focus On" section.
- **Chapter 4** illustrates basic computer organization and introduces many fundamental concepts, including the fetch–decode–execute cycle, the data path, clocks and buses, register transfer notation, and, of course, the CPU. A very simple architecture, MARIE, and its ISA are presented to allow the reader to gain a full understanding of the basic architectural organization involved in program execution. MARIE exhibits the classic von Neumann design and includes a program counter, an accumulator, an instruction register, 4096 bytes of memory, and two addressing modes. Assembly language programming is introduced to reinforce the concepts of instruction format, instruction mode, data format, and control that are presented earlier. This is not an assembly language textbook and was not designed to provide a practical course in assembly language programming. The primary objective in introducing assembly is to further the understanding of computer architecture in general. However, a simulator for MARIE is provided so assembly language programs can be written, assembled, and run on the MARIE architecture. The two methods of control, hardwiring and microprogramming, are introduced and compared in this chapter. Finally, Intel and MIPS architectures are compared to reinforce the concepts in the chapter.
- **Chapter 5** provides a closer look at instruction set architectures, including instruction formats, instruction types, and addressing modes. Instruction-level pipelining is introduced as well. Real-world ISAs (including Intel®, MIPS® Technologies, ARM, and Java™) are presented to reinforce the concepts presented in the chapter.

- **Chapter 6** covers basic memory concepts, such as RAM and the various memory devices, and also addresses the more advanced concepts of the memory hierarchy, including cache memory and virtual memory. This chapter gives a thorough presentation of direct mapping, associative mapping, and set-associative mapping techniques for cache. It also provides a detailed look at paging and segmentation, TLBs, and the various algorithms and devices associated with each. A tutorial and simulator for this chapter is available on the book's website.
- **Chapter 7** provides a detailed overview of I/O fundamentals, bus communication and protocols, and typical external storage devices, such as magnetic and optical disks, as well as the various formats available for each. DMA, programmed I/O, and interrupts are covered as well. In addition, various techniques for exchanging information between devices are introduced. RAID architectures are covered in detail. Various data compression formats are introduced in a special "Focus On" section.
- **Chapter 8** discusses the various programming tools available (such as compilers and assemblers) and their relationship to the architecture of the machine on which they are run. The goal of this chapter is to tie the programmer's view of a computer system with the actual hardware and architecture of the underlying machine. In addition, operating systems are introduced, but only covered in as much detail as applies to the architecture and organization of a system (such as resource use and protection, traps and interrupts, and various other services).
- **Chapter 9** provides an overview of alternative architectures that have emerged in recent years. RISC, Flynn's Taxonomy, parallel processors, instruction-level parallelism, multiprocessors, interconnection networks, shared memory systems, cache coherence, memory models, superscalar machines, neural networks, systolic architectures, dataflow computers, quantum computing, and distributed architectures are covered. Our main objective in this chapter is to help the reader realize we are not limited to the von Neumann architecture, and to force the reader to consider performance issues, setting the stage for the next chapter.
- **Chapter 10** covers concepts and topics of interest in embedded systems that have not been covered in previous chapters. Specifically, this chapter focuses on embedded hardware and components, embedded system design topics, the basics of embedded software construction, and embedded operating systems features.
- **Chapter 11** addresses various performance analysis and management issues. The necessary mathematical preliminaries are introduced, followed by a discussion of MIPS, FLOPS, benchmarking, and various optimization issues with which a computer scientist should be familiar, including branch prediction, speculative execution, and loop optimization.
- **Chapter 12** focuses on network organization and architecture, including network components and protocols. The OSI model and TCP/IP suite are introduced in the context of the Internet. This chapter is by no means intended to

be comprehensive. The main objective is to put computer architecture in the correct context relative to network architecture. Chapter 12 is now available online.

- **Chapter 13** introduces some popular I/O architectures suitable for large and small systems, including SCSI, ATA, IDE, SATA, PCI, USB, and IEEE 1394. This chapter also provides a brief overview of storage area networks and cloud computing. Chapter 13 is now available online.
- **Appendix A** is a short appendix on data structures that is provided for those situations in which students may need a brief introduction or review of such topics as stacks, queues, and linked lists. Appendix A is now available online.

The sequencing of the chapters is such that they can be taught in the given numerical order. However, an instructor can modify the order to better fit a given




FIGURE P.1 Prerequisite Relationship Between Chapters

curriculum if necessary. Figure P.1 shows the prerequisite relationships that exist between various chapters.

What's New in the Fifth Edition

In the years since the fourth edition of this book was created, the field of computer architecture has continued to grow. In this fifth edition, we have incorporated many of these changes in addition to expanding topics already introduced in the first four editions. Our goal in the fifth edition was to update content in all chapters, add new material, and expand current discussions based on reader comments. In addition, we have added sidebars, called “Null Pointers,” throughout the textbook to focus on potential pitfalls and provide helpful hints and tips. Some content (see following) has been moved online in an effort to streamline the textbook. Although we cannot itemize all the changes in this edition, the list that follows highlights those major changes that may be of interest to the reader:

- **Chapter 1** has been updated to include new examples, as well as additional information on smart phones and tablets. The hardware overview has been updated, and further information has been added regarding input/output components, the ALU, memory, the data path, and interconnection structures (memory modules, I/O modules, the CPU, and buses). Transistors have been deemphasized and more importance has been placed on computer systems. The storage hierarchy is introduced, and various types of memory (such as RAM and RDRAM) are covered in more detail. The discussion of cloud computing has been expanded and the basic concepts have been made more accessible. Pertinent sections of Chapter 12 have been integrated into Chapter 1.
- **Chapter 2** has dropped the author-invented format for floating point and now introduces the floating-point format using the IEEE format. More details and examples have been provided for floating-point numbers, and the “Codes for Data Recording and Transmission” section has been moved online.
- **Chapter 3** brings Kmaps back into the main chapter, devotes more time to optimization, and includes the concepts of prime implicants and essential implicants. More arithmetic circuits have been introduced, and step-by-step techniques for analyzing and designing flip-flops and synchronous sequential machines have been added. A section on hardware description languages (including Verilog and VHDL) has been included, and the section on finite state machines has been rewritten to make it more accessible and to relate it to concepts introduced in Chapter 4.
- **Chapter 4** has expanded sections to include more MIPS and Intel programs for comparison with MARIE programs. In addition, there are now more exercises for the MARIE assembly language.
- **Chapter 5** has been updated to include the MIPS architecture in the “real world section.” It also has additional material to help make design choices more relevant by connecting the architectural considerations to the implications these decisions have on the programmer.

- **Chapter 6** expands the real-world memory study and includes a new section on the MIPS architecture.
- **Chapter 7** has an updated disk technology section, and data compression has been moved online. The discussion of how interrupts work has been expanded.
- **Chapter 8** has been updated by moving the database software and transaction manager sections online. More focus has been placed on software/hardware interaction.
- **Chapter 12** (what remains after integrating parts into Chapter 1) has been moved online.
- **Chapter 13** has been moved online.
- **Appendix A** has been moved online.

For more information on accessing online material, please visit go.jblearning.com/null5e or contact your Account Representative at go.jblearning.com/findmyrep.

Intended Audience

This book was originally written for an undergraduate class in computer organization and architecture for computer science majors. Although specifically directed toward computer science majors, the book does not preclude its use by IS and IT majors.

This book contains more than sufficient material for a typical one-semester (14 weeks, 42 lecture hours) course; however, all the material in the book cannot be mastered by the average student in a one-semester class. If the instructor plans to cover all topics in detail, a two-semester sequence would be optimal.

Chapter	One Semester (42 Hours)		Two Semesters (84 Hours)	
	Lecture Hours	Expected Level	Lecture Hours	Expected Level
1	3	Mastery	3	Mastery
2	6	Mastery	6	Mastery
3	6	Mastery	6	Mastery
4	8	Mastery	8	Mastery
5	4	Familiarity	6	Mastery
6	3	Familiarity	8	Mastery
7	2	Familiarity	6	Mastery
8	2	Exposure	7	Mastery
9	2	Familiarity	7	Mastery
10	1	Exposure	5	Familiarity
11	2	Exposure	9	Mastery
12	2	Exposure	7	Mastery
13	1	Exposure	6	Mastery

TABLE P.2 Suggested Lecture Hours

The organization is such that an instructor can cover the major topic areas at different levels of depth, depending on the experience and needs of the students. Table P.2 gives the instructor an idea of the amount of time required to cover the topics, and also lists the corresponding levels of accomplishment for each chapter.

It is our intention that this book serve as a useful reference long after the formal course is complete.

Support Materials

A textbook is a fundamental tool in learning, but its effectiveness is greatly enhanced by supplemental materials and exercises, which emphasize the major concepts, provide immediate feedback to the reader, and motivate understanding through repetition. We have, therefore, created the following ancillary materials for the fifth edition of *The Essentials of Computer Organization and Architecture*. For more information on how to access these resources, please visit go.jblearning.com/null5e.

- *Test Bank*.
- *Instructor's Manual*. This manual contains answers to exercises. In addition, it provides hints on teaching various concepts and trouble areas often encountered by students.
- *PowerPoint Presentations*. These slides contain lecture material appropriate for a one-semester course in computer organization and architecture.
- *Figures and Tables*. For those who wish to prepare their own lecture materials, we provide the figures and tables in downloadable form.
- *Memory Tutorial and Simulator*. This package allows students to apply the concepts on cache and virtual memory.
- *MARIE Simulator*. This package allows students to assemble and run MARIE programs.
- *Datapath Simulator*. This package allows students to trace the MARIE datapath.
- *Tutorial Software*. Other tutorial software is provided for various concepts in the book.
- *An Introduction to MIPS Assembly Language*. This supplementary eBook on MIPS Assembly Language can be purchased separately or bundled with the textbook.
- *Intel Assembly Language*. This supplement provides additional information on the Intel Assembly Language.

The exercises, sample exam problems, and solutions have been tested in numerous classes. The *Instructor's Manual*, which includes suggestions for teaching the various chapters in addition to answers for the book's exercises, suggested programming assignments, and sample example questions, is available to instructors who adopt the book. (Please contact your Jones & Bartlett Learning representative at 1-800-832-0034 or visit go.jblearning.com/findmyrep for access to this area of the website.)

The Instructional Model: MARIE

In a computer organization and architecture book, the choice of architectural model affects the instructor as well as the students. If the model is too complicated, both the instructor and the students tend to get bogged down in details that really have no bearing on the concepts being presented in class. Real architectures, although interesting, often have far too many peculiarities to make them usable in an introductory class. To make things even more complicated, real architectures change from day to day. In addition, it is difficult to find a book incorporating a model that matches the local computing platform in a given department, noting that the platform, too, may change from year to year.

To alleviate these problems, we have designed our own simple architecture, MARIE, specifically for pedagogical use. MARIE (Machine Architecture that is Really Intuitive and Easy) allows students to learn the essential concepts of computer organization and architecture, including assembly language, without getting caught up in the unnecessary and confusing details that exist in real architectures. Despite its simplicity, it simulates a functional system. The MARIE machine simulator, MarieSim, has a user-friendly GUI that allows students to (1) create and edit source code, (2) assemble source code into machine object code, (3) run machine code, and (4) debug programs.

Specifically, MarieSim has the following features:

- Support for the MARIE assembly language introduced in Chapter 4
- An integrated text editor for program creation and modification
- Hexadecimal machine language object code
- An integrated debugger with single step mode, break points, pause, resume, and register and memory tracing
- A graphical memory monitor displaying the 4096 addresses in MARIE's memory
- A graphical display of MARIE's registers
- Highlighted instructions during program execution
- User-controlled execution speed
- Status messages
- User-viewable symbol tables
- An interactive assembler that lets the user correct any errors and reassemble automatically, without changing environments
- Online help
- Optional core dumps, allowing the user to specify the memory range
- Frames with sizes that can be modified by the user
- A small learning curve, allowing students to learn the system quickly

MarieSim was written in the Java language so that the system would be portable to any platform for which a Java Virtual Machine (JVM) is available. Students of Java may wish to look at the simulator's source code, and perhaps even offer improvements or enhancements to its simple functions.




FIGURE P.2 The MarieSim Graphical Environment

Figure P.2, the MarieSim Graphical Environment, shows the graphical environment of the MARIE machine simulator. The screen consists of four parts: the menu bar, the central monitor area, the memory monitor, and the message area.

Menu options allow the user to control the actions and behavior of the MARIE simulator system. These options include loading, starting, stopping, setting breakpoints, and pausing programs that have been written in MARIE assembly language.

The MARIE simulator illustrates the process of assembly, loading, and execution, all in one simple environment. Users can see assembly language statements directly from their programs, along with the corresponding machine code (hexadecimal) equivalents. The addresses of these instructions are indicated as well, and users can view any portion of memory at any time. Highlighting is used to indicate the initial loading address of a program in addition to the currently executing instruction while a program runs. The graphical display of the registers and memory allows the student to see how the instructions cause the values in the registers and memory to change.

If You Find an Error

We have attempted to make this book as technically accurate as possible, but even though the manuscript has been through numerous proofreadings, errors have a way of escaping detection. We would greatly appreciate hearing from readers who find any errors that need correcting. Your comments and suggestions are always welcome; please send an email to ECOIA@jblearning.com.

Credits and Acknowledgments

Few books are entirely the result of one or two people's unaided efforts, and this one is no exception. We realize that writing a textbook is a formidable task and only possible with a combined effort, and we find it impossible to adequately thank those who have made this book possible. If, in the following acknowledgments, we inadvertently omit anyone, we humbly apologize.

A number of people have contributed to the fifth edition of this book. We would first like to thank all of the reviewers for their careful evaluations of previous editions and their thoughtful written comments. In addition, we are grateful for the many readers who have emailed useful ideas and helpful suggestions. We extend a special thanks to Karishma Rao and Sean Willeford for their time and effort in producing a quality memory software module.

We would also like to thank the individuals at Jones & Bartlett Learning who worked closely with us to make this fifth edition possible. We are very grateful for their professionalism, commitment, and hard work on the fifth edition.

I, Linda Null, would personally like to posthumously thank my husband, Tim Wahls, for his patience and for putting up with the almost daily compromises necessitated by my writing this book—including missing our annual fly-fishing vacation and forcing our horses into prolonged pasture ornament status. I consider myself amazingly lucky to have been married to such a wonderful man and will miss him terribly. I extend my heartfelt thanks to my mentor, Merry McDonald, who taught me the value and joys of learning and teaching, and doing both with integrity. I would also like to thank my friend and colleague, Thang Bui, for all of his support and encouragement from the very beginning. Lastly, I would like to express my deepest gratitude to Julia Lobur, as without her, this book and its accompanying software would not be a reality. It has been both a joy and an honor working with her.

I, Julia Lobur, am deeply indebted to my lawful spouse, Marla Cattermole, who married me despite the demands that this book has placed on both of us. She has made this work possible through her forbearance and fidelity. She has nurtured my body through her culinary delights and my spirit through her wisdom. She has taken up my slack in many ways while working hard at her own career. I would also like to convey my profound gratitude to Linda Null: first, for her unsurpassed devotion to the field of computer science education and dedication to her students and, second, for giving me the opportunity to share with her the inefable experience of textbook authorship.

