

# 《计算机体系结构》

## 图书基本信息

书名：《计算机体系结构》

13位ISBN编号：9787111364580

10位ISBN编号：7111364589

出版时间：2012-1

出版社：机械工业出版社

作者：John L. Hennessy, David A. Patterson

页数：856

版权说明：本站所提供下载的PDF图书仅提供预览和简介以及在线试读，请支持正版图书。

更多资源请访问：[www.tushu000.com](http://www.tushu000.com)

# 《计算机体系结构》

## 内容概要

# 《计算机体系结构》

## 作者简介

John L. Hennessy 斯坦福大学校长，IEEE和ACM会士，美国国家工程研究院院士及美国科学艺术研究院院士。Hennessy教授因为在RISC技术方面做出了突出贡献而荣获2001年的Eckert-Mauchly奖章，他也是2001年Seymour Cray计算机工程奖得主，并且和本书另外一位作者David A. Patterson分享了2000年John von Neumann奖。

David A. Patterson 加州大学伯克利分校计算机科学系主任、教授，美国国家工程研究院院士，IEEE和ACM会士，曾因成功的启发式教育方法被IEEE授予James H. Mulligan, Jr.教育奖章。他因为对RISC技术的贡献而荣获1995年IEEE技术成就奖，而在RAID技术方面的成就为他赢得了1999年IEEE Reynold Johnson信息存储奖。2000年他和John L. Hennessy分享了John von Neumann奖。

## 书籍目录

Foreword

Preface

Acknowledgments

Chapter 1 Fundamentals of Quantitative Design and Analysis

1.1 Introduction

1.2 Classes of Computers

1.3 Defining Computer Architecture

1.4 Trends in Technology

1.5 Trends in Power and Energy in Integrated Circuits

1.6 Trends in Cost

1.7 Dependability

1.8 Measuring, Reporting, and Summarizing Performance

1.9 Quantitative Principles of Computer Design

1.10 Putting It All Together: Performance, Price, and Power

1.11 Fallacies and Pitfalls

1.12 Concluding Remarks

1.13 Historical Perspectives and References Case Studies and Exercises by Diana Franklin

Chapter 2 Memory Hierarchy Design

2.1 Introduction

2.2 Ten Advanced Optimizations of Cache Performance

2.3 Memory Technology and Optimizations

2.4 Protection: Virtual Memory and Virtual Machines

2.5 Crosscutting Issues: The Design of Memory Hierarchies

2.6 Putting It All Together: Memory Hierarchies in the ARM Cortex-A5 and Intel Core i7

2.7 Fallacies and Pitfalls

2.8 Concluding Remarks: Looking Ahead

2.9 Historical Perspective and References Case Studies and Exercises by Norman P. Jouppi, Naveen Muralimanohar, and Sheng Li

Chapter 3 Instruction-Level Parallelism and Its Exploitation

3.1 Instruction-Level Parallelism: Concepts and Challenges

3.2 Basic Compiler Techniques for Exposing ILP

3.3 Reducing Branch Costs with Advanced Branch Prediction

3.4 Overcoming Data Hazards with Dynamic Scheduling

3.5 Dynamic Scheduling: Examples and the Algorithm

3.6 Hardware-Based Speculation

3.7 Exploiting ILP Using Multiple Issue and Static Scheduling

3.8 Exploiting ILP Using Dynamic Scheduling, Multiple Issue, and Speculation

3.9 Advanced Techniques for Instruction Delivery and Speculation

3.10 Studies of the Limitations of ILP

3.11 Cross-Cutting Issues: ILP Approaches and the Memory System

3.12 Multithreading: Exploiting Thread-Level Parallelism to Improve Uniprocessor Throughput

3.13 Putting It All Together: The Intel Core i7 and ARM Cortex-A5

3.14 Fallacies and Pitfalls

3.15 Concluding Remarks: What's Ahead?

3.16 Historical Perspective and References Case Studies and Exercises by Jason D. Bakos and Robert R Colwell

Chapter 4 Data-Level Parallelism in Vector, SIMD, and GPU Architectures

4.1 Introduction

- 4.2 Vector Architecture
- 4.3 SIMD Instruction Set Extensions for Multimedia
- 4.4 Graphics Processing Units
- 4.5 Detecting and Enhancing Loop-Level Parallelism
- 4.6 Crosscutting Issues
- 4.7 Putting It All Together: Mobile versus Server GPUS and Tesla versus Core i7
- 4.8 Fallacies and Pitfalls
- 4.9 Concluding Remarks
- 4.10 Historical Perspective and References Case Study and Exercises by Jason D. Bakos
- Chapter 5 Thread-Level Parallelism
- 5.1 Introduction
- 5.2 Centralized Shared-Memory Architectures
- 5.3 Performance of Symmetric Shared-Memory Multiprocessors
- 5.4 Distributed Shared-Memory and Directory-Based Coherence
- 5.5 Synchronization: The Basics
- 5.6 Models of Memory Consistency: An Introduction
- 5.7 Crosscutting Issues
- 5.8 Putting It All Together: Multicore Processors and Their Performance
- 5.9 Fallacies and Pitfalls
- 5.10 Concluding Remarks
- 5.11 Historical Perspectives and References
- Case Studies and Exercises by Amr Zaky and David A. Wood
- Chapter 6 Warehouse-Scale Computers to Exploit Request-Level and Data-Level Parallelism
- 6.1 Introduction
- 6.2 Programming Models and Workloads for Warehouse-Scale Computers
- 6.3 Computer Architecture of Warehouse-Scale Computers
- 6.4 Physical Infrastructure and Costs of Warehouse-Scale Computers
- 6.5 Cloud Computing: The Return of Utility Computing
- 6.6 Crosscutting Issues
- 6.7 Putting It All Together: A Google Warehouse-Scale Computer
- 6.8 Fallacies and Pitfalls
- 6.9 Concluding Remarks
- 6.10 Historical Perspectives and References
- Case Studies and Exercises by Parthasarathy Ranganathan
- Appendix A Instruction Set Principles
- A.1 Introduction
- A.2 Classifying Instruction Set Architectures
- A.3 Memory Addressing
- A.4 Type and Size of Operands
- A.5 Operations in the Instruction Set
- A.6 Instructions for Control Flow
- A.7 Encoding an Instruction Set
- A.8 Crosscutting Issues: The Role of Compilers
- A.9 Putting It All Together: The MIPS Architecture
- A.10 Fallacies and Pitfalls
- A.11 Concluding Remarks
- A.12 Historical Perspective and References
- Exercises by Gregory D. Peterson

## Appendix B Review of Memory Hierarchy

### B.1 Introduction

### B.2 Cache Performance

### B.3 Six Basic Cache Optimizations

### B.4 Virtual Memory

### B.5 Protection and Examples of Virtual Memory

### B.6 Fallacies and Pitfalls

### B.7 Concluding Remarks

### B.8 Historical Perspective and References

### Exercises by Amr Zaky

## Appendix C Pipelining: Basic and Intermediate Concepts

### C.1 Introduction

### C.2 The Major Hurdle of Pipelining--Pipeline Hazards

### C.3 How Is Pipelining Implemented?

### C.4 What Makes Pipelining Hard to Implement?

### C.5 Extending the MIPS Pipeline to Handle Multicycle Operations

### C.6 Putting It All Together: The MIPS R4000 Pipeline

### C.7 Crosscutting Issues

### C.8 Fallacies and Pitfalls

### C.9 Concluding Remarks

### C.10 Historical Perspective and References

### Updated Exercises by Diana Franklin

### Online Appendices

## Appendix D Storage Systems

## Appendix E Embedded Systems

### By ThomasM Conte

## Appendix F Interconnection Networks

### Revised by Timothy M. Pinkston and Jose Duoto

## Appendix G Vector Processors in More Depth

### Revised by Krste Asanovic

## Appendix H Hardware and Software for VLIW and EPIC

## Appendix I Large-Scale Multiprocessors and Scientific Applications

## Appendix J Computer Arithmetic

### by David Goldberg

## Appendix K Survey of Instruction Set Architectures

## Appendix L Historical Perspectives and References

## References

## Index

### 1.1 Introduction

### 1.2 Classes of Computers

### 1.3 Defining Computer Architecture

### 1.4 Trends in Technology

### 1.5 Trends in Power and Energy in Integrated Circuits

### 1.6 Trends in Cost

### 1.7 Dependability

### 1.8 Measuring Reporting, and Summarizing Performance

### 1.9 Quantitative Principles of Computer Design

### 1.10 Putting It All Together: Performance, Price, and Power

### 1.11 Fallacies and Pitfalls

1.12 Concluding Remarks

1.13 Historical Perspectives and References

1.14 Case Studies and Exercises by Diana Franklin

## 章节摘录

版权页：插图：The pressure of both a fast clock cycle and power limitations encourages limited size for first-level caches. Similarly, use of lower levels of associativity can reduce both hit time and power, although such trade-offs are more complex than those involving size. The critical timing path in a cache hit is the three-step process of addressing the tag memory using the index portion of the address, comparing the read tag value to the address, and setting the multiplexor to choose the correct data item if the cache is set associative. Direct-mapped caches can overlap the tag check with the transmission of the data, effectively reducing hit time. Furthermore, lower levels of associativity will usually reduce power because fewer cache lines must be accessed. Although the total amount of on-chip cache has increased dramatically with new generations of microprocessors, due to the clock rate impact arising from a larger L1 cache, the size of the L1 caches has recently increased either slightly or not at all. In many recent processors, designers have opted for more associativity rather than larger caches. An additional consideration in choosing the associativity is the possibility of eliminating address aliases; we discuss this shortly. One approach to determining the impact on hit time and power consumption in advance of building a chip is to use CAD tools. CACTI is a program to estimate the access time and energy consumption of alternative cache structures on CMOS microprocessors within 10% of more detailed CAD tools. For a given minimum feature size, CACTI estimates the hit time of caches as cache size varies, associativity, number of read/write ports, and more complex parameters. Figure 2.3 shows the estimated impact on hit time as cache size and associativity are varied.



# 《计算机体系结构》

## 媒体关注与评论

“ 本书之所以成为永恒的经典，是因为它的每一次再版都不仅仅是更新补充，而是一次全面的修订，对这个激动人心且快速变化领域给出了最及时的信息和最独到的解读。对于我来说，即使已有二十多年的从业经历，再次阅读本书仍自觉学无止境，感佩于两位卓越大师的渊博学识和深厚功底。” ——Luiz Andre Barroso，Google公司



## 精彩短评

### 1、体系结构研究入门书

我的英语书籍阅读的第一个啃的书籍

2、中英文的间杂着看，中文的太多瑕疵了，勘误了一堆。看了一部分只有，应试。

3、趁活动买的，经典书籍，就不多评论了

4、每个CS人都应该读的书籍。

5、还没看，不过书的包装和质量还是感觉不错的。希望内容能有一些惊喜

6、厚厚一本，经典的书籍，讲的很清晰

7、重新开读。

8、哎，感觉不适合做教材

9、书的印刷还好，书里的插图丰富，每页都留有足够空白位置做笔记，但为什么不像CSAPP那样双色印刷呢？就算贵一点也没关系

10、体系结构经典，买过好几版，内容不断在与时俱进。

11、买的这本书，价格不算便宜。可是书的质量实在不敢恭维。里面有倒页，而且不止一处。我向亚马逊反映，也没有给我满意的答复。如果说书的印刷装订质量有问题，亚马逊没办法保证，那是印刷厂出版社的责任。但如果再有卖家的服务问题，那真是无语了。建议大家买前考虑清楚。

12、《软硬件接口》进阶版，紧贴时代潮流。

13、每天多看你一眼，了解就多一分。第一次见识loop unrolling是在programming pearls上面，然后就很多次见到这个词，这本书给出了更多解释。随着Microsoft, google, Facebook等大量集群系统出现，他们不采用传统数据库系统，而是自行发展自己的软件平台，最后一章对这些进行了精彩的阐述。至于其余:ILP 本书一直以来保留项目,存储器层次结构 保留项目,TLP 多处理器时代的主题,DLP 混合架构CPU+GPU嗯，期待能更深刻地理解计算机的工作方式

14、两个作者：一个是斯坦福校长，创建了MIPS；一个是伯克利计算机系主任，开发了RAID。经典不解释！

15、这本书改变了我对计算机的认识

16、书很不错，我们需要的就是这个英文版的

17、深入浅出。别忘了看附录。

18、经典书籍，正在研读中

19、计算机体系结构的圣经。

20、经典著作，无需多言。

21、很厚实，里面看了下还不错，就是教科书类型

22、当初学的死去活来。

23、这本书浅显易懂。块头很大，是因为语言组织的很容易理解！想深入理解计算机结构的人可以好好钻研下这本书！

24、经典书籍，值得看看

25、这一版的笔误明显比上一版多，让人有些失望。新增加的WSC一章，对开拓学生的视野很有好处，不过把集群看做一台计算机，个人观点，略显牵强了。

26、相对第4版：更新了每章后面的实例分析部分，增加了GPU相关章节，总的来说改动不大

27、体系结构领域的圣经，还有什么好说呢。

28、经典教材，体系结构专业必备

29、虽然上了这门课，用的这本教材，真不敢说读了这本书。

30、大开眼界

31、13/11/22

因为蛋疼的243开始读的.发现写的确实很赞~

32、英文原版计算机教材，对计算机体系结构有很好的阐述

33、手头的工具书，对于计算机硬件的讲解真的非常棒，可惜处于时间缘故，仅仅挑着读了其中几章，目前流行的GPU缺没有读，非常遗憾。

34、希望有用处。。。。。。。。。。

## 《计算机体系结构》

- 35、教材。。。
- 36、想不到竟然这么好，获益良多，应该可以反复读，喜欢大部头
- 37、还是第三版好些
- 38、研究生的第一学期彻底被这本书毁了。尼玛第五版的习题要不要这么难！！！！
- 39、经典之作，值得慢慢读
- 40、内容很好，就是关键词看起来吃力一些
- 41、看了一章，感觉很不错。不过亚马逊发的书有些脏，机械工业出版社的标签被撕掉一半，看起来不像新书。
- 42、计算机体系结构方面的经典之作
- 43、Nice book. Clear statement, variety of examples.
- 44、北大CS研究生教材
- 45、原来hennessy写的是这本...书没看懂但今年也算是现场听过他扯淡了
- 46、第5版！
- 47、绝对是计算机体系结构方面的圣经啊，书的质量没有原版的好，但是还是不错了，比原版便宜了不少。

# 《计算机体系结构》

## 精彩书评

1、北大东门附近交易吧。第五版，英文版，机械工业出版的，9成新，原价138，80出。电话：  
133411267三七啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦啦  
啦啦啦啦啦啦啦啦啦啦

# 《计算机体系结构》

## 版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:[www.tushu000.com](http://www.tushu000.com)