Innovations in the Memory System

Synthesis Lectures on Computer Architecture

Editors
Natalie Enright Jerger, ***University of Toronto***
Margaret Martonosi, ***Princeton University***

Founding Editor Emeritus
Mark D. Hill, ***University of Wisconsin, Madison***

Synthesis Lectures on Computer Architecture publishes 50- to 100-page publications on topics pertaining to the science and art of designing, analyzing, selecting and interconnecting hardware components to create computers that meet functional, performance and cost goals. The scope will largely follow the purview of premier computer architecture conferences, such as ISCA, HPCA, MICRO, and ASPLOS.

Innovations in the Memory System
Rajeev Balasubramonian
2019

Cache Replacement Policies
Akanksha Jain and Calvin Lin
2019

The Datacenter as a Computer: Designing Warehouw-Scale Machines, Third Edition
Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan
2018

Principles of Secure Processor Architecture Design
Jakub Szefer
2018

General-Purpose Graphics Processor Architectures
Tor M. Aamodt, Wilson Wai Lun Fung, and Timothy G. Rogers
2018

Compiling Algorithms for Heterogenous Systems
Steven Bell, Jing Pu, James Hegarty, and Mark Horowitz
2018

Architectural and Operating System Support for Virtual Memory
Abhishek Bhattacharjee and Daniel Lustig
2017

Deep Learning for Computer Architects
Brandon Reagen, Robert Adolf, Paul Whatmough, Gu-Yeon Wei, and David Brooks
2017

On-Chip Networks, Second Edition
Natalie Enright Jerger, Tushar Krishna, and Li-Shiuan Peh
2017

Space-Time Computing with Temporal Neural Networks
James E. Smith
2017

Hardware and Software Support for Virtualization
Edouard Bugnion, Jason Nieh, and Dan Tsafrir
2017

Datacenter Design and Management: A Computer Architect's Perspective
Benjamin C. Lee
2016

A Primer on Compression in the Memory Hierarchy
Somayeh Sardashti, Angelos Arelakis, Per Stenström, and David A. Wood
2015

Research Infrastructures for Hardware Accelerators
Yakun Sophia Shao and David Brooks
2015

Analyzing Analytics
Rajesh Bordawekar, Bob Blainey, and Ruchir Puri
2015

Customizable Computing
Yu-Ting Chen, Jason Cong, Michael Gill, Glenn Reinman, and Bingjun Xiao
2015

Die-stacking Architecture
Yuan Xie and Jishen Zhao
2015

Single-Instruction Multiple-Data Execution
Christopher J. Hughes
2015

Power-Efficient Computer Architectures: Recent Advances
Magnus Själander, Margaret Martonosi, and Stefanos Kaxiras
2014

FPGA-Accelerated Simulation of Computer Systems
Hari Angepat, Derek Chiou, Eric S. Chung, and James C. Hoe
2014

A Primer on Hardware Prefetching
Babak Falsafi and Thomas F. Wenisch
2014

On-Chip Photonic Interconnects: A Computer Architect's Perspective
Christopher J. Nitta, Matthew K. Farrens, and Venkatesh Akella
2013

Optimization and Mathematical Modeling in Computer Architecture
Tony Nowatzki, Michael Ferris, Karthikeyan Sankaralingam, Cristian Estan, Nilay Vaish, and David Wood
2013

Security Basics for Computer Architects
Ruby B. Lee
2013

The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Second Edition
Luiz André Barroso, Jimmy Clidaras, and Urs Hölzle
2013

Shared-Memory Synchronization
Michael L. Scott
2013

Resilient Architecture Design for Voltage Variation
Vijay Janapa Reddi and Meeta Sharma Gupta
2013

Multithreading Architecture
Mario Nemirovsky and Dean M. Tullsen
2013

Performance Analysis and Tuning for General Purpose Graphics Processing Units (GPGPU)
Hyesoon Kim, Richard Vuduc, Sara Baghsorkhi, Jee Choi, and Wen-mei Hwu
2012

Automatic Parallelization: An Overview of Fundamental Compiler Techniques
Samuel P. Midkiff
2012

Phase Change Memory: From Devices to Systems
Moinuddin K. Qureshi, Sudhanva Gurumurthi, and Bipin Rajendran
2011

Multi-Core Cache Hierarchies
Rajeev Balasubramonian, Norman P. Jouppi, and Naveen Muralimanohar
2011

A Primer on Memory Consistency and Cache Coherence
Daniel J. Sorin, Mark D. Hill, and David A. Wood
2011

Dynamic Binary Modification: Tools, Techniques, and Applications
Kim Hazelwood
2011

Quantum Computing for Computer Architects, Second Edition
Tzvetan S. Metodi, Arvin I. Faruque, and Frederic T. Chong
2011

High Performance Datacenter Networks: Architectures, Algorithms, and Opportunities
Dennis Abts and John Kim
2011

Processor Microarchitecture: An Implementation Perspective
Antonio González, Fernando Latorre, and Grigorios Magklis
2010

Transactional Memory, Second Edition
Tim Harris, James Larus, and Ravi Rajwar
2010

Computer Architecture Performance Evaluation Methods
Lieven Eeckhout
2010

Introduction to Reconfigurable Supercomputing
Marco Lanzagorta, Stephen Bique, and Robert Rosenberg
2009

On-Chip Networks
Natalie Enright Jerger and Li-Shiuan Peh
2009

The Memory System: You Can't Avoid It, You Can't Ignore It, You Can't Fake It
Bruce Jacob
2009

Fault Tolerant Computer Architecture
Daniel J. Sorin
2009

The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines
Luiz André Barroso and Urs Hölzle
2009

Computer Architecture Techniques for Power-Efficiency
Stefanos Kaxiras and Margaret Martonosi
2008

Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency
Kunle Olukotun, Lance Hammond, and James Laudon
2007

Transactional Memory
James R. Larus and Ravi Rajwar
2006

Quantum Computing for Computer Architects
Tzvetan S. Metodi and Frederic T. Chong
2006

© Springer Nature Switzerland AG 2022
Reprint of original edition © Morgan & Claypool 2019

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Innovations in the Memory System

Rajeev Balasubramonian

ISBN: 978-3-031-00635-7 paperback
ISBN: 978-3-031-01763-6 ebook
ISBN: 978-3-031-00060-7 hardcover

DOI 10.1007/978-3-031-01763-6

A Publication in the Springer series
SYNTHESIS LECTURES ON COMPUTER ARCHITECTURE

Lecture #48
Series Editors: Natalie Enright Jerger, *University of Toronto*
Margaret Martonosi, *Princeton University*
Founding Editor Emeritus: Mark D. Hill, *University of Wisconsin, Madison*
Series ISSN
Print 1935-3235 Electronic 1935-3243

Innovations in the Memory System

Rajeev Balasubramonian
University of Utah

SYNTHESIS LECTURES ON COMPUTER ARCHITECTURE #48

ABSTRACT

The memory system has the potential to be a hub for future innovation. While conventional memory systems focused primarily on high density, other memory system metrics like energy, security, and reliability are grabbing modern research headlines. With processor performance stagnating, it is also time to consider new programming models that move some application computations into the memory system. This, in turn, will lead to feature-rich memory systems with new interfaces. The past decade has seen a number of memory system innovations that point to this future where the memory system will be much more than dense rows of unintelligent bits. This book takes a tour through recent and prominent research works, touching upon new DRAM chip designs and technologies, near data processing approaches, new memory channel architectures, techniques to tolerate the overheads of refresh and fault tolerance, security attacks and mitigations, and memory scheduling.

KEYWORDS

memory system architecture, DRAM, memory controllers, memory channels, low-power memory, memory security, error correction codes, memory reliability, DRAM refresh, near data processing, 3D stacking

Contents

List of Figures xv

List of Tables xvii

Preface xix

Acknowledgments xxi

1 Introduction 1

2 Memory System Basics for Every Architect 3

2.1 DRAM vs. SRAM 3
2.2 Memory Channel 4
2.3 DDR Standards 5
2.4 DIMMs, Ranks, Banks, Mats 5
2.5 Row Buffers 7
2.6 Capacity vs. Energy 8
2.7 Address Mapping 9
2.8 Scheduling 9
2.9 DRAM Timing Parameters 10

3 Commercial Memory Products 13

3.1 Basic DDR3/DDR4 Channels and DIMMs 13
3.2 DDR Deviations for Higher Capacity and Bandwidth 16

4 Memory Scheduling 21

4.1 Memory Scheduler Basics 21
4.2 Early Multi-Core Memory Schedulers 22
4.3 Co-Designed Schedulers 24
4.4 Discussion 26

5 Data Placement . . . 27
5.1 Data Interleaving . . . 27
5.2 Memory Compression . . . 28
5.2.1 IBM MXT . . . 28
5.2.2 Ekman and Stenstrom . . . 29
5.2.3 Linearly Compressed Pages (LCP) . . . 29
5.2.4 Nearly Overhead-Free Memory Compression . . . 30
5.2.5 PTMC . . . 32
5.2.6 Active Memory Expansion in the IBM Power Processors . . . 32
5.3 Discussion . . . 33

6 Memory Chip Microarchitectures . . . 35
6.1 Basics of DRAM Chip Microarchitecture . . . 35
6.2 DRAM Chip Innovations . . . 37
6.3 Discussion . . . 40

7 Memory Channels . . . 41
7.1 The Basics of Parallel and Serial Interconnects . . . 41
7.1.1 Parallel Buses . . . 41
7.1.2 Serial Buses . . . 42
7.1.3 Discussion . . . 43
7.2 Memory Interconnect Innovations . . . 45
7.2.1 Recent Work . . . 45
7.2.2 Discussion . . . 49

8 Memory Reliability . . . 51
8.1 Basics of DRAM Errors . . . 51
8.2 Memory Reliability Innovations . . . 54
8.3 Discussion . . . 58

9 Memory Refresh . . . 59
9.1 Refresh Basics . . . 59
9.2 Refresh Innovations . . . 63
9.2.1 Empirical Studies and Retention Times . . . 63
9.2.2 Alternative Hardware Techniques . . . 65
9.2.3 Alternative Software Techniques . . . 66
9.2.4 Leveraging Charge in Cells . . . 66
9.3 Discussion . . . 67

10 Near Data Processing . . . 69
10.1 NDP Implementations . . . 70
10.1.1 3D Stacked Architectures . . . 70
10.1.2 Tight Coupling on a DIMM . . . 71
10.2 In-Situ Implementations . . . 73
10.3 Programming Models and Applications . . . 77
10.3.1 Programming Approaches . . . 77
10.3.2 Broadening the Scope of NDP with Additional Workloads . . . 79
10.4 Discussion . . . 80

11 Memory Security . . . 81
11.1 Memory Timing Channels . . . 81
11.2 Oblivious RAM (ORAM) . . . 87
11.3 Memory Integrity . . . 91
11.4 Impact of Smart Memories . . . 98
11.5 Other Memory Security Issues . . . 99
11.6 Discussion . . . 100

12 Closing Thoughts . . . 103

Bibliography . . . 105

Author's Biography . . . 129

List of Figures

2.1 A comparison of SRAM and DRAM cell circuits 4
2.2 DRAM hierarchy of DIMM, rank, chip, bank, subarray, mat 6

3.1 Example processor/memory configuration 14
3.2 Organization of buffered DIMMs 15
3.3 FB-DIMM and Buffer-on-Board organizations 17
3.4 The Hybrid Memory Cube organization 19

6.1 DRAM array microarchitecture 36

7.1 Memory power for HMC-based organizations 44
7.2 Memory I/O power for DDR3 and DDR4 45

10.1 Dot-product computation with a memristor crossbar 74

11.1 Performance for timing-channel-free memory controllers 85
11.2 Timing-channel-free memory controller scheduling policies 86
11.3 Path ORAM example 89
11.4 Merkle Trees 94
11.5 Intel SGX integrity tree 95
11.6 VAULT integrity tree 96

List of Tables

2.1 DRAM timing parameter descriptions 11
2.1 DRAM timing parameter descriptions 12

3.1 Summary of DRAM parameters across generations 16

9.1 Refresh latencies for high DRAM chip capacities 63

Preface

The memory sub-system is the most overlooked and underrated component in modern systems. Memory chips are just rows of dumb bits: they are commoditized and standardized; they perform two simple operations—reads and writes.

A lot of the above conventional wisdom is changing. With processor performance improving only slightly with each new generation, the computing industry needs creative value-additions in new systems. Computation is widely viewed as "trivial" or "commoditized," while data movement is viewed as the engineer's nemesis. Most solutions to these problems seem to rely on two broad compelling approaches: accelerators and new memory technologies.

For decades, architecture conferences saw little activity in memory systems. In the past ten years, this area has been a dominant topic, only recently ceding that title to machine learning accelerators. Much of this recent excitement stemmed from new non-volatile memory technologies and 3D stacking. Both made it possible to improve energy efficiency and add more functionality to the memory system. Once the floodgates opened, researchers have delved into several aspects of memory systems, including security, reliability, refresh, microarchitecture, etc.

However, even today, most computer architecture textbooks spend hundreds of pages on the processor and just a handful of pages on the memory system. Most graduate students starting their Ph.D.s in architecture require a solid foundation in memory design. We hope that this Synthesis Lecture serves as a valuable resource to such students, as well as to other architecture researchers/developers. Instructors teaching an advanced computer architecture class may want to base some of their lectures on parts of this book. Chapter 2 would also be appropriate reading for a standard course on computer architecture. To make the material easily accessible to early-stage graduate students, each chapter starts by covering the basics and then diving into key ideas/contributions in recent top-tier papers. The goal here is to provide a condensed big-picture overview of the landscape; once informed, students can chase down specific topics/papers that deserve more attention. While I have read hundreds of papers so you don't have to :-), I highly recommend that you absolutely read the original papers that you find relevant. I also want to recommend the 2008 memory system book written by Jacob, Ng, and Wang [4]; it's a terrific reference that I have frequently consulted.

Given the many popular conferences/workshops in existence today, I have obviously not covered every published memory system innovation. To be fair and realistic, I have primarily covered papers published at the top four architecture venues—ISCA, MICRO, ASPLOS, HPCA. I realize that this approach has shortcomings, and I apologize in advance if I did not mention your favorite work.

The book represents the hard work of a dedicated team at Morgan & Claypool, the careful oversight of series editors Margaret Martonosi and Natalie Enright Jerger, the constructive feedback of reviewers, and the education I have received from my students and collaborators. We hope you enjoy the book!

Rajeev Balasubramonian
July 2019

Acknowledgments

I am very grateful to all my collaborators, especially Al Davis and all my Ph.D. students who have taught me so much about each aspect of the memory system. I can claim to be a memory expert only because my students' intellectual curiosity, creativity, and attention to detail pulled my research program in new and exciting directions. As always, a ton of credit goes to my parents and my highly supportive family—Deepthi, Shurik, and Anushka.

Rajeev Balasubramonian
July 2019