Datacenter Design and Management

A Computer Architect's Perspective

Synthesis Lectures on Computer Architecture

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, 2nd 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

Datacenter Design and Management: A Computer Architect's Perspective

Benjamin C. Lee

ISBN: 978-3-031-00624-1 paperback
ISBN: 978-3-031-01752-0 ebook

DOI 10.1007/978-3-031-01752-0

A Publication in the Springer series
SYNTHESIS LECTURES ON COMPUTER ARCHITECTURE

Lecture #37
Series Editor: Margaret Martonosi, *Princeton University*
Series ISSN
Print 1935-3235 Electronic 1935-3243

Datacenter Design and Management

A Computer Architect's Perspective

Benjamin C. Lee
Duke University

SYNTHESIS LECTURES ON COMPUTER ARCHITECTURE #37

ABSTRACT

An era of big data demands datacenters, which house the computing infrastructure that translates raw data into valuable information. This book defines datacenters broadly, as large distributed systems that perform parallel computation for diverse users. These systems exist in multiple forms—private and public—and are built at multiple scales. Datacenter design and management is multifaceted, requiring the simultaneous pursuit of multiple objectives. Performance, efficiency, and fairness are first-order design and management objectives, which can each be viewed from several perspectives. This book surveys datacenter research from a computer architect's perspective, addressing challenges in applications, design, management, server simulation, and system simulation. This perspective complements the rich bodies of work in datacenters as a warehouse-scale system, which study the implications for infrastructure that encloses computing equipment, and in datacenters as distributed systems, which employ abstract details in processor and memory subsystems. This book is written for first- or second-year graduate students in computer architecture and may be helpful for those in computer systems. The goal of this book is to prepare computer architects for datacenter-oriented research by describing prevalent perspectives and the state-of-the-art.

KEYWORDS

computer organization and design, energy efficiency, cluster computing, data centers, distributed systems, cloud computing, performance evaluation methodologies, resource allocation, software scheduling

Contents

Preface

This synthesis lecture is written for first- or second-year graduate students in computer architecture. The reader is expected to have completed graduate coursework in computer architecture; additionally, a course in distributed systems would be helpful. Moreover, the reader is expected to have some basic knowledge and prior experience in using the tools of the trade: cycle-level simulators. This background provides the requisite perspective on benchmarking and simulation for conventional workloads to help the reader appreciate challenges that are new and unique to datacenter workloads.

Moreover, this lecture may be helpful for graduate students in computer systems. Because the determinants of datacenter performance and efficiency increasingly lie at the hardware-software interface, architecture and systems perspectives on datacenters could be integrated to reveal new research directions. Because datacenter operators are rightfully wary of introducing new hardware into well-tuned systems, architects must anticipate system management challenges during architectural design. Furthermore, emerging hardware technologies and architectures require new system organizations and management.

Related Lectures. This synthesis lecture complements two existing synthesis lectures, one on datacenters and another on performance evaluation methods. Both lectures are highly recommended for their breadth and complementary perspective. Barroso et al. present a lecture on datacenters that focuses on the design of warehouse-scale machines, which is often taken to mean the datacenter itself [16]. For example, the lecture describes the facility, the peripheral infrastructure that supports the computing equipment, and figures of merit for evaluating datacenter efficiency and costs (e.g., total cost of ownership). This lecture is highly recommended for its breadth, its focus on warehouse-scale systems and facilities, and its industry-strength perspective. In contrast, our lecture focuses on processor and memory design, and emphasizes experimental methodologies that draw on a rich body of widely deployed open-source applications.

Eeckhout's lecture on performance evaluation methods focuses on performance evaluation methodologies, with a specific emphasis on strategies that accelerate the evaluation process [42]. The lecture describes analytical performance models that concisely represent processor performance. It also describes varied statistical strategies that reveal application performance while reducing the number of instructions simulated with cycle-level timing models. These research methodologies are best suited for understanding broad, general-purpose benchmark suites. In contrast, we focus on datacenter workloads and full system simulation. Previously proposed strategies for rapid design space exploration may apply to datacenter workloads as well, but we would need to adapt them to full system simulation.

We organize this lecture on datacenter research methodologies in several chapters. Chapter 2 describes several representative datacenter applications, surveys their implications for hardware architectures, and proposes benchmarking strategies. Chapters 3–4 survey recent research in server design and management. Chapter 5 details strategies for simulating datacenter servers. Specifically, we present approaches to processor and memory simulation, and demonstrate methodologies for precise simulations that target application regions of interest. Finally, Chapter 6 describes strategies for simulating datacenter dynamics at scale. We present analytical and empirical approaches to understanding task behaviors and queueing dynamics.

Collectively, the goal of this book is to prepare a computer architect for datacenter-oriented research. It describes prevalent perspectives and the state-of-the-art. Yet, for all the research that is surveyed in this book, many challenges remain and the required advances in datacenter design and management would very much benefit from a computer architect's perspective.

Benjamin C. Lee
January 2016

Acknowledgments

This synthesis lecture draws from varied research projects that have spanned the last eight years. During this time, I have been fortunate to work with extraordinary collaborators and students. My perspectives on computer architecture and datacenter systems have been enriched by these collaborations. I am thankful for my collaborators during projects at Microsoft Research—Vijay Janapa Reddi, Trishul Chilimbi, and Kushagra Vaid—an institution where fundamental research is leavened with applied perspectives. I am also grateful for my collaborators during projects at Stanford University—Mark Horowitz, Christos Kozyrakis, and Krishna Malladi—a place where interdisciplinary work is pervasive. Finally, I am incredibly fortunate to work with a group of outstanding Ph.D. students at Duke University—Songchun Fan, Marisabel Guevara, Ziqiang Huang, Tamara Lehman, Bryan Prosser, Qiuyun Wang, Seyed Majid Zahedi, and Pengfei Zheng.

For this project, I am thankful for Margaret Martonosi and Michael Morgan's invitation to write a synthesis lecture on datacenter research. The lecture has benefited substantially from reviews by Christina Delimitrou, Stijn Eyerman, Boris Grot, and Lingjia Tang. In advance, I thank the faculty and students who use and share this lecture to advance their own research projects. I would very much appreciate any feedback from researchers and practitioners.

Lastly, I would like to thank our research sponsors for their support of my projects and the students who contribute to them. Our projects have been supported by the National Science Foundation, Duke University, and the Google Faculty Research Award. They have also been supported by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA. Any opinions, findings, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of these sponsors.

Benjamin C. Lee
January 2016