Analytical Performance Modeling for Computer Systems

Third Edition

Synthesis Lectures on Computer Science

Analytical Performance Modeling for Computer Systems, Third Edition
Y.C. Tay
2018

Creating Autonomous Vehicle Systems
Shaoshan Liu, Liyun Li, Jie Tang, Shuang Wu, and Jean-Luc Gaudiot
2017

Introduction to Logic, Third Edition
Michael Genesereth and Eric Kao
2016

Analytical Performance Modeling for Computer Systems, Second Edition
Y.C. Tay
2013

Introduction to Logic, Second Edition
Michael Genesereth and Eric Kao
2013

Introduction to Logic
Michael Genesereth and Eric Kao
2013

Science Fiction Prototyping: Designing the Future with Science Fiction
Brian David Johnson
2011

Storing Clocked Programs Inside DNA: A Simplifying Framework for Nanocomputing
Jessica P. Chang and Dennis E. Shasha
2011

Analytical Performance Modeling for Computer Systems
Y.C. Tay
2010

The Theory of Timed I/O Automata
Dilsun K. Kaynar, Nancy Lynch, Roberto Segala, and Frits Vaandrager
2006

Analytical Performance Modeling for Computer Systems, Third Edition

Y.C. Tay

ISBN: 978-3-031-00675-3 paperback
ISBN: 978-3-031-01803-9 ebook
ISBN: 978-3-031-00071-3 hardcover

DOI 10.1007/978-3-031-01803-9

A Publication in the Springer series
SYNTHESIS LECTURES ON COMPUTER SCIENCE

Lecture #10
Series ISSN
Synthesis Lectures on Computer Science
Print 1932-1228 Electronic 1932-1686

Analytical Performance Modeling for Computer Systems

Third Edition

Y.C. Tay
National University of Singapore

SYNTHESIS LECTURES ON COMPUTER SCIENCE #10

ABSTRACT

This book is an introduction to analytical performance modeling for computer systems, i.e., writing equations to describe their performance behavior. It is accessible to readers who have taken college-level courses in calculus and probability, networking, and operating systems.

This is not a training manual for becoming an expert performance analyst. Rather, the objective is to help the reader construct simple models for analyzing and understanding the systems that they are interested in.

Describing a complicated system abstractly with mathematical equations requires a careful choice of assumptions and approximations. They make the model tractable, but they must not remove essential characteristics of the system, nor introduce spurious properties.

To help the reader understand the choices and their implications, this book discusses the analytical models for 40 research papers. These papers cover a broad range of topics: GPUs and disks, routers and crawling, databases and multimedia, worms and wireless, multicore and cloud, security and energy, etc. An appendix provides many questions for readers to exercise their understanding of the models in these papers.

KEYWORDS

computer system performance, analytical modeling techniques, simulation, experimental validation, Markov chains, queueing systems, fluid approximation, transient analysis

Contents

Preface

In writing this book, I have in mind the student, engineer or researcher who

(a) is interested in the performance of some particular computer system, and
(b) wants to analytically model that behavior, but
(c) does not intend to become an expert in performance analysis.

For network systems, the literature has numerous examples of analytical performance models; evidently, this research community has found such models useful. This book would have served its purpose if it can help the researchers in other communities (hardware architecture, operating systems, programming languages, database management, etc.) add a modeling chapter to a thesis or a similar section in a paper.

There is a common perception that performance modeling requires a lot of queueing theory. This is not so. Queueing systems are used in this book only as an expository device for a coherent presentation. The concepts (e.g., open/closed in Chapter 1, residual life in Chapter 3, flow equivalence in Chapter 6, stability in Chapter 8), techniques (e.g., Markov chains in Chapter 4, Average Value Approximation and fluid approximation in Chapter 7, equilibrium decomposition in Chapter 8), and results (Little's Law in Chapter 1, effect of variance on response time in Chapter 2, PASTA in Chapter 5, bottleneck analysis in Chapter 6) are applicable to more than just queueing systems.

In the first edition, I chose 20 papers to illustrate the ideas in the book. The papers were as follows.

(1) *StreamJoins*

J. Kang, J. F. Naughton, and S. Viglas. Evaluating window joins over unbounded streams. In *Proc. IEEE Int. Conf. on Data Engineering (ICDE)*, 341–352, March 2003.

(2) *SleepingDisks*

Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes. Hibernator: helping disk arrays sleep through the winter. In *Proc. ACM Symp. Operating Systems Principles (SOSP)*, 39(5):177–190, October 2005.

(3) *GPRS*

G. Nogueira, B. Baynat, and P. Eisenmann. An analytical model for the dimensioning of a GPRS/EDGE network with a capacity constraint on a group of cells. In *Proc. ACM MOBICOM*, 215–227, August 2005.

(4) *InternetServices*

B. Urgaonkar, G. Pacifici, P. Shenoy, M. Spreitzer, and A. Tantawi. Analytic modeling of multitier Internet applications. *ACM Trans. Web*, 1(1):2, 2007.

(5) *OpenClosed*

B. Schroeder, A. Wierman, and M. Harchol-Balter. Open versus closed: a cautionary tale. In *Proc. Symp. Networked Systems Design and Implementation (NSDI)*, May 2006.

(6) *TCP*

J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP throughput: a simple model and its empirical validation. In *Proc. ACM SIGCOMM*, 303–314, September 1998.

(7) *BitTorrent*

D. Qiu and R. Srikant. Modeling and performance analysis of BitTorrent-like peer-to-peer networks. In *Proc. ACM SIGCOMM*, 367–378, 2004.

(8) *CodeRed*

C. C. Zou, W. Gong, and D. Towsley. Code red worm propagation modeling and analysis. In *Proc. ACM Conf. Computer and Communications Security (CCS)*, 138–147, November 2002.

(9) *WirelessCapacity*

P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, 46(2):388–404, 2000.

(10) *802.11*

Y. C. Tay and K. C. Chua. A capacity analysis for the IEEE 802.11 MAC protocol. *Wireless Networks*, 7(2):159–171, 2001.

(11) *MediaStreaming*

Y.-C. Tu, J. Sun, and S. Prabhakar. Performance analysis of a hybrid media streaming system. In *Proc. ACM/SPIE Conf. on Multimedia Computing and Networking (MMCN)*, 69–82, January 2004.

(12) *StorageAvailability*

E. Gabber, J. Fellin, M. Flaster, F. Gu, B. Hillyer, W. T. Ng, B. Özden, and E. A. M. Shriver. Starfish: highly-available block storage. In *Proc. USENIX Annual Tech. Conf.*, 151–163, June 2003.

(13) *TransactionalMemory*

A. Heindl, G. Pokam, and A.-R. Adl-Tabatabai. An analytic model of optimistic software transactional memory. In *Proc. IEEE Int. Symp. on Performance Analysis of Systems and Software (ISPASS)*, 153–162, April 2009.

(14) *SensorNet*
R. C. Shah, S. Roy, S. Jain, and W. Brunette. Data mules: Modeling a three-tier architecture for sparse sensor networks. In *Proc. IEEE Workshop on Sensor Network Protocols and Applications*, 30–41, May 2003.

(15) *NetworkProcessor*
J. Lu and J. Wang. Analytical performance analysis of network processor-based application design. In *Proc. Int. Conf. Computer Communications and Networks*, 78–86, October 2006.

(16) *DatabaseSerializability*
P. A. Bernstein, A. Fekete, H. Guo, R. Ramakrishnan, and P. Tamma. Relaxed currency serializability for middle-tier caching and replication. In *Proc. ACM SIGMOD Int. Conf. Management of Data*, 599–610, June 2006.

(17) *NoC*
J. Kim, D. Park, C. Nicopoulos, N. Vijaykrishnan, and C. R. Das. Design and analysis of an NoC architecture from performance, reliability and energy perspective. In *Proc. ACM Symp. Architecture for Networking and Communications Systems (ANCS)*, 173–182, October 2005.

(18) *DistributedProtocols*
I. Gupta. On the design of distributed protocols from differential equations. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, 216–225, July 2004.

(19) *NonstationaryMix*
C. Stewart, T. Kelly, and A. Zhang. Exploiting nonstationarity for performance prediction. In *Proc. ACM EuroSys Conf.*, 31–44, March 2007.

(20) *SoftState*
J. C. S. Lui, V. Misra, and D. Rubenstein. On the robustness of soft state protocols. In *Proc. IEEE Int. Conf. Network Protocols (ICNP)*, 50–60, October 2004.

In the second edition, I added 10 papers (and exercises) to illustrate the concepts, techniques, and results in this book, and demonstrate breadth in their application. Several were chosen to show how an analytical model can offer insight not available through experiments. These 10 papers were as follows.

(21) *WebCrawler*
J. Cho and H. Garcia-Molina. The evolution of the web and implications for an incremental crawler. In *Proc. Int. Conf. on Very Large Data Bases (VLDB)*, 200–209, September 2000.

(22) *DatacenterAMP*
V. Gupta and R. Nathuji. Analyzing performance asymmetric multicore processors for

latency sensitive datacenter applications. In *Proc. Int. Conf. Power Aware Computing and Systems*, 1–8, October 2010.

(23) *RouterBuffer*
G. Appenzeller, I. Keslassy, and N. McKeown. Sizing router buffers. In *Proc. ACM SIGCOMM*, 281–292, August 2004.

(24) *DependabilitySecurity*
K. S. Trivedi, D. S. Kim, A. Roy, and D. Medhi. Dependability and security models. In *Proc. Int. Workshop on Design of Reliable Communication Networks (DRCN)*, 11–20, October 2009.

(25) *PipelineParallelism*
A. Navarro, R. Asenjo, Si. Tabik, and C. Cascaval. Analytical modeling of pipeline parallelism. In *Proc. Int. Conf. on Parallel and Architectures and Compilation Techniques (PACT)*, 281–290, September 2009.

(26) *Roofline*
S. Williams, A. Waterman, and D. Patterson. Roofline: an insightful visual performance model for multicore architectures. *Commun. ACM*, 52(4):65–76, April 2009.

(27) *Gossip*
R. Bakhshi, D. Gavidia, W. Fokkink, and M. Steen. An analytical model of information dissemination for a gossip-based protocol. In *Proc. Int. Conf. Distributed Computing and Networking (ICDCN)*, 230–242, January 2009.

(28) *EpidemicRouting*
X. Zhang, G. Neglia, J. Kurose, and D. Towsley. Performance modeling of epidemic routing. *Computer Networks*, 51(10):2867–2891, July 2007.

(29) *P2PVoD*
B. Ran, D. G. Andersen, M. Kaminsky, and K. Papagiannaki. Balancing throughput, robustness, and in-order delivery in P2P VoD. In *Proc. CoNEXT*, 10:1–10:12, November 2010.

(30) *CloudTransactions*
D. Kossman, T. Kraska, and S. Loesing. An evaluation of alternative architectures for transaction processing in the cloud. In *Proc. ACM SIGMOD Int. Conf. Management of Data*, 579–590, June 2010.

For the third edition, I used similar criteria to select another 10 papers (with corresponding exercises).

(31) *SoftErrors*
A.A. Nair, S. Eyerman, L. Eeckhout, and L.K. John. A first-order mechanistic model for architectural vulnerability factor. In *Proc. IEEE Int. Symp. Computer Architecture (ISCA)*, 273–284, June 2012.

(32) *GPU*
J.-C. Huang, J.H. Lee, H. Kim, and H.-H. S. Lee. GPUMech: GPU performance modeling technique based on interval analysis. In *Proc. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, 268–279, December 2014.

(33) *ProactiveReplication*
A. Duminuco, E. Biersack, and T. En-Najjary. Proactive replication in distributed storage systems using machine availability Estimation. In *Proc. ACM CoNEXT Conf.*, 27:1–27:12, December 2007.

(34) *ServerEnergy*
B. Guenter, N. Jain, and C. Williams. Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning. In *Proc. IEEE INFOCOM*, 1332–1340, April 2011.

(35) *DatabaseScalability*
S. Elnikety, S. Dropsho, E. Cecchet, and W. Zwaenepoel. Predicting replicated database scalability from standalone database profiling. In *Proc. ACM EuroSys Conf.*, 303–316, April 2009.

(36) *PerformanceAssurance*
N. Roy, A. Dubey, A. Gokhale, and L. Dowdy. A capacity planning process for performance assurance of component-based distributed Systems. In *Proc. ACM/SPEC Int. Conf. Performance Engineering (ICPE)*, 259–270, September 2011.

(37) *CachingSystems*
V. Martina, M. Garetto, and E. Leonardi. A unified approach to the performance analysis of caching systems. In *Proc. IEEE INFOCOM*, 2040–2048, April 2014.

(38) *InformationDiffusion*
Y. Matsubara, Y. Sakurai, B.A. Prakash, L. Li and C. Faloutsos. Rise and fall patterns of information diffusion: model and implications. In *Proc. ACM KDD*, 6–14, August 2012.

(39) *MapReduce*
E. Vianna, G. Comarela, T. Pontes, J. Almeida, V. Almeida, K. Wilkinson, H. Kuno, and U. Dayal. Analytical performance models for MapReduce workloads. *Int. J. Parallel Programming* 41(4):495–525, August 2013.

(40) *ElasticScaling*
D. Didona, P. Romano, S. Peluso, and F. Quaglia. Transactional Auto Scaler: Elastic scaling of in-memory transactional data grids. In *Proc. Int. Conf. Autonomic Computing (ICAC)*, 125–134, September 2012.

These papers were also chosen to demonstrate various aspects of performance modeling: applying the MVA algorithm to nonseparable queueing networks (*DatabaseScalability* and *PerformanceAssurance*), how to factor in the effect of precedence constraints on delays (*MapReduce*), unrolling a state transition diagram (*ServerEnergy*), how a model can be used to control a system (*ProactiveReplication*), the difference between whitebox and blackbox models (*SoftErrors* and *ElasticScaling*), how machine learning can help (*GPU* and *ElasticScaling*), and using a model to develop the science in system behavior (*CachingSystems* and *InformationDiffusion*).

But most important of all, I hope the readers who I have in mind can gain confidence from studying these examples, and see that they too can construct an analytical model for the performance behavior of their system.

Y.C. Tay
Kent Ridge
Westwood 2010
Palo Alto 2013
Taipei 2017
June 2018

Acknowledgments

The drafts of this book were tested with courses at NUS (National University of Singapore), UCLA (University of California, Los Angeles) and NTU (National Taiwan University). Many thanks to my hosts Dick Muntz at UCLA and Yuh-Dauh Lyuu at NTU, and to the reviewers for their encouraging comments. I finished the second edition while on sabbatical at VMware; I thank the many colleagues there who made my visit so enjoyable and productive. Thanks too to Diane Cerra of Morgan & Claypool for her encouragement, and Clovis L. Tondo for his invaluable help with the typesetting.

Y.C. Tay
June 2018