

1. Record Nr.	UNINA9910465041203321
Autore	Harchol-Balter Mor <1966->
Titolo	Performance modeling and design of computer systems : queueing theory in action // Mor Harchol-Balter, Carnegie Mellon University, Pennsylvania [[electronic resource]]
Pubbl/distr/stampa	Cambridge : , : Cambridge University Press, , 2013
ISBN	1-107-23684-3 1-107-25475-2 1-139-61641-2 1-139-62571-3 1-139-61269-7 1-139-22642-8
Descrizione fisica	1 online resource (xxiii, 548 pages) : digital, PDF file(s)
Disciplina	519.8/2
Soggetti	Transaction systems (Computer systems) - Mathematical models Computer systems - Design and construction - Mathematics Queueing theory Queueing networks (Data transmission)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Title from publisher's bibliographic system (viewed on 05 Oct 2015).
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Machine generated contents note: Part I. Introduction to Queueing: 1. Motivating examples; 2. Queueing theory terminology; Part II. Necessary Probability Background: 3. Probability review; 4. Generating random variables; 5. Sample paths, convergence, and averages; Part III. The Predictive Power of Simple Operational Laws: 'What If' Questions and Answers; 6. Operational laws; 7. Modification analysis; Part IV. From Markov Chains to Simple Queues: 8. Discrete-time Markov Chains; 9. Ergodicity theory; 10. Real-world examples: Google, Aloha; 11. Generating functions for Markov Chains; 12. Exponential distributions and Poisson Process; 13. Transition to continuous-time Markov Chains; 14. M/M/1 and PASTA; Part V. Server Farms and Networks: Multi-server, Multi-queue Systems: 15. Server farms: M/M/k and M/M/k/k; 16. Capacity provisioning for server farms; 17. Time-reversibility and Burke's Theorem; 18. Jackson network of queues; 19.

Classed network of queues; 20. Closed networks of queues; Part VI. Real-World Workloads: High-Variability and Heavy Tails: 21. Tales of tails: real-world workloads; 22. Phase-type workloads and matrix-analytic; 23. Networks of time-sharing (PS) servers; 24. M/G/I queue and inspection paradox; 25. Task assignment for server farms; 26. Transform analysis; 27. M/G/I transform analysis; 28. Power optimization application; Part VII. Smart Scheduling: 29. Performance metrics; 30. Non-preemptive, non-size-based policies; 31. Preemptive, non-size-based policies; 32. Non-preemptive, size-based policies; 33. Preemptive, size-based policies; 34. Scheduling: SRPT and fairness.

Sommario/riassunto

Tackling the questions that systems designers care about, this book brings queueing theory decisively back to computer science. The book is written with computer scientists and engineers in mind and is full of examples from computer systems, as well as manufacturing and operations research. Fun and readable, the book is highly approachable, even for undergraduates, while still being thoroughly rigorous and also covering a much wider span of topics than many queueing books. Readers benefit from a lively mix of motivation and intuition, with illustrations, examples and more than 300 exercises - all while acquiring the skills needed to model, analyze and design large-scale systems with good performance and low cost. The exercises are an important feature, teaching research-level counterintuitive lessons in the design of computer systems. The goal is to train readers not only to customize existing analyses but also to invent their own.
