

Contents

<i>Preface</i>	xvii
<i>Acknowledgments</i>	xxiii

I Introduction to Queueing

1 Motivating Examples of the Power of Analytical Modeling	3
1.1 What Is Queueing Theory?	3
1.2 Examples of the Power of Queueing Theory	5
2 Queueing Theory Terminology	13
2.1 Where We Are Heading	13
2.2 The Single-Server Network	13
2.3 Classification of Queueing Networks	16
2.4 Open Networks	16
2.5 More Metrics: Throughput and Utilization	17
2.6 Closed Networks	20
2.6.1 Interactive (Terminal-Driven) Systems	21
2.6.2 Batch Systems	22
2.6.3 Throughput in a Closed System	23
2.7 Differences between Closed and Open Networks	24
2.7.1 A Question on Modeling	25
2.8 Related Readings	25
2.9 Exercises	26

II Necessary Probability Background

3 Probability Review	31
3.1 Sample Space and Events	31
3.2 Probability Defined on Events	32
3.3 Conditional Probabilities on Events	33
3.4 Independent Events and Conditionally Independent Events	34
3.5 Law of Total Probability	35
3.6 Bayes Law	36
3.7 Discrete versus Continuous Random Variables	37
3.8 Probabilities and Densities	38
3.8.1 Discrete: Probability Mass Function	38
3.8.2 Continuous: Probability Density Function	41
3.9 Expectation and Variance	44
3.10 Joint Probabilities and Independence	47

viii	CONTENTS	
3.11	Conditional Probabilities and Expectations	49
3.12	Probabilities and Expectations via Conditioning	53
3.13	Linearity of Expectation	54
3.14	Normal Distribution	57
3.14.1	Linear Transformation Property	58
3.14.2	Central Limit Theorem	61
3.15	Sum of a Random Number of Random Variables	62
3.16	Exercises	64
4	Generating Random Variables for Simulation	70
4.1	Inverse-Transform Method	70
4.1.1	The Continuous Case	70
4.1.2	The Discrete Case	72
4.2	Accept-Reject Method	72
4.2.1	Discrete Case	73
4.2.2	Continuous Case	75
4.2.3	Some Harder Problems	77
4.3	Readings	78
4.4	Exercises	78
5	Sample Paths, Convergence, and Averages	79
5.1	Convergence	79
5.2	Strong and Weak Laws of Large Numbers	83
5.3	Time Average versus Ensemble Average	84
5.3.1	Motivation	85
5.3.2	Definition	86
5.3.3	Interpretation	86
5.3.4	Equivalence	88
5.3.5	Simulation	90
5.3.6	Average Time in System	90
5.4	Related Readings	91
5.5	Exercise	91
III The Predictive Power of Simple Operational Laws: “What-If” Questions and Answers		
6	Little’s Law and Other Operational Laws	95
6.1	Little’s Law for Open Systems	95
6.2	Intuitions	96
6.3	Little’s Law for Closed Systems	96
6.4	Proof of Little’s Law for Open Systems	97
6.4.1	Statement via Time Averages	97
6.4.2	Proof	98
6.4.3	Corollaries	100
6.5	Proof of Little’s Law for Closed Systems	101
6.5.1	Statement via Time Averages	101
6.5.2	Proof	102
6.6	Generalized Little’s Law	102

CONTENTS	ix
6.7 Examples Applying Little's Law	103
6.8 More Operational Laws: The Forced Flow Law	106
6.9 Combining Operational Laws	107
6.10 Device Demands	110
6.11 Readings and Further Topics Related to Little's Law	111
6.12 Exercises	111
7 Modification Analysis: "What-If" for Closed Systems	114
7.1 Review	114
7.2 Asymptotic Bounds for Closed Systems	115
7.3 Modification Analysis for Closed Systems	118
7.4 More Modification Analysis Examples	119
7.5 Comparison of Closed and Open Networks	122
7.6 Readings	122
7.7 Exercises	122
IV From Markov Chains to Simple Queues	
8 Discrete-Time Markov Chains	129
8.1 Discrete-Time versus Continuous-Time Markov Chains	130
8.2 Definition of a DTMC	130
8.3 Examples of Finite-State DTMCs	131
8.3.1 Repair Facility Problem	131
8.3.2 Umbrella Problem	132
8.3.3 Program Analysis Problem	132
8.4 Powers of \mathbf{P} : n -Step Transition Probabilities	133
8.5 Stationary Equations	135
8.6 The Stationary Distribution Equals the Limiting Distribution	136
8.7 Examples of Solving Stationary Equations	138
8.7.1 Repair Facility Problem with Cost	138
8.7.2 Umbrella Problem	139
8.8 Infinite-State DTMCs	139
8.9 Infinite-State Stationarity Result	140
8.10 Solving Stationary Equations in Infinite-State DTMCs	142
8.11 Exercises	145
9 Ergodicity Theory	148
9.1 Ergodicity Questions	148
9.2 Finite-State DTMCs	149
9.2.1 Existence of the Limiting Distribution	149
9.2.2 Mean Time between Visits to a State	153
9.2.3 Time Averages	155
9.3 Infinite-State Markov Chains	155
9.3.1 Recurrent versus Transient	156
9.3.2 Infinite Random Walk Example	160
9.3.3 Positive Recurrent versus Null Recurrent	162
9.4 Ergodic Theorem of Markov Chains	164

x	CONTENTS	
9.5	Time Averages	166
9.6	Limiting Probabilities Interpreted as Rates	168
9.7	Time-Reversibility Theorem	170
9.8	When Chains Are Periodic or Not Irreducible	171
9.8.1	Periodic Chains	171
9.8.2	Chains that Are Not Irreducible	177
9.9	Conclusion	177
9.10	Proof of Ergodic Theorem of Markov Chains*	178
9.11	Exercises	183
10	Real-World Examples: Google, Aloha, and Harder Chains*	190
10.1	Google's PageRank Algorithm	190
10.1.1	Google's DTMC Algorithm	190
10.1.2	Problems with Real Web Graphs	192
10.1.3	Google's Solution to Dead Ends and Spider Traps	194
10.1.4	Evaluation of the PageRank Algorithm	195
10.1.5	Practical Implementation Considerations	195
10.2	Aloha Protocol Analysis	195
10.2.1	The Slotted Aloha Protocol	196
10.2.2	The Aloha Markov Chain	196
10.2.3	Properties of the Aloha Markov Chain	198
10.2.4	Improving the Aloha Protocol	199
10.3	Generating Functions for Harder Markov Chains	200
10.3.1	The z-Transform	201
10.3.2	Solving the Chain	201
10.4	Readings and Summary	203
10.5	Exercises	204
11	Exponential Distribution and the Poisson Process	206
11.1	Definition of the Exponential Distribution	206
11.2	Memoryless Property of the Exponential	207
11.3	Relating Exponential to Geometric via δ -Steps	209
11.4	More Properties of the Exponential	211
11.5	The Celebrated Poisson Process	213
11.6	Merging Independent Poisson Processes	218
11.7	Poisson Splitting	218
11.8	Uniformity	221
11.9	Exercises	222
12	Transition to Continuous-Time Markov Chains	225
12.1	Defining CTMCs	225
12.2	Solving CTMCs	229
12.3	Generalization and Interpretation	232
12.3.1	Interpreting the Balance Equations for the CTMC	234
12.3.2	Summary Theorem for CTMCs	234
12.4	Exercises	234

CONTENTS	xi
13 M/M/1 and PASTA	236
13.1 The M/M/1 Queue	236
13.2 Examples Using an M/M/1 Queue	239
13.3 PASTA	242
13.4 Further Reading	245
13.5 Exercises	245
V Server Farms and Networks: Multi-server, Multi-queue Systems	
14 Server Farms: M/M/k and M/M/k/k	253
14.1 Time-Reversibility for CTMCs	253
14.2 M/M/k/k Loss System	255
14.3 M/M/k	258
14.4 Comparison of Three Server Organizations	263
14.5 Readings	264
14.6 Exercises	264
15 Capacity Provisioning for Server Farms	269
15.1 What Does Load Really Mean in an M/M/k?	269
15.2 The M/M/∞	271
15.2.1 Analysis of the M/M/∞	271
15.2.2 A First Cut at a Capacity Provisioning Rule for the M/M/k	272
15.3 Square-Root Staffing	274
15.4 Readings	276
15.5 Exercises	276
16 Time-Reversibility and Burke's Theorem	282
16.1 More Examples of Finite-State CTMCs	282
16.1.1 Networks with Finite Buffer Space	282
16.1.2 Batch System with M/M/2 I/O	284
16.2 The Reverse Chain	285
16.3 Burke's Theorem	288
16.4 An Alternative (Partial) Proof of Burke's Theorem	290
16.5 Application: Tandem Servers	291
16.6 General Acyclic Networks with Probabilistic Routing	293
16.7 Readings	294
16.8 Exercises	294
17 Networks of Queues and Jackson Product Form	297
17.1 Jackson Network Definition	297
17.2 The Arrival Process into Each Server	298
17.3 Solving the Jackson Network	300
17.4 The Local Balance Approach	301
17.5 Readings	306
17.6 Exercises	306
18 Classed Network of Queues	311
18.1 Overview	311
18.2 Motivation for Classed Networks	311

xii	CONTENTS	
18.3	Notation and Modeling for Classed Jackson Networks	314
18.4	A Single-Server Classed Network	315
18.5	Product Form Theorems	317
18.6	Examples Using Classed Networks	322
	18.6.1 Connection-Oriented ATM Network Example	322
	18.6.2 Distribution of Job Classes Example	325
	18.6.3 CPU-Bound and I/O-Bound Jobs Example	326
18.7	Readings	329
18.8	Exercises	329
19	Closed Networks of Queues	331
19.1	Motivation	331
19.2	Product-Form Solution	333
	19.2.1 Local Balance Equations for Closed Networks	333
	19.2.2 Example of Deriving Limiting Probabilities	335
19.3	Mean Value Analysis (MVA)	337
	19.3.1 The Arrival Theorem	338
	19.3.2 Iterative Derivation of Mean Response Time	340
	19.3.3 An MVA Example	341
19.4	Readings	343
19.5	Exercises	343
 VI Real-World Workloads: High Variability and Heavy Tails 		
20	Tales of Tails: A Case Study of Real-World Workloads	349
20.1	Grad School Tales . . . Process Migration	349
20.2	UNIX Process Lifetime Measurements	350
20.3	Properties of the Pareto Distribution	352
20.4	The Bounded Pareto Distribution	353
20.5	Heavy Tails	354
20.6	The Benefits of Active Process Migration	354
20.7	Pareto Distributions Are Everywhere	355
20.8	Exercises	357
21	Phase-Type Distributions and Matrix-Analytic Methods	359
21.1	Representing General Distributions by Exponentials	359
21.2	Markov Chain Modeling of PH Workloads	364
21.3	The Matrix-Analytic Method	366
21.4	Analysis of Time-Varying Load	367
	21.4.1 High-Level Ideas	367
	21.4.2 The Generator Matrix, Q	368
	21.4.3 Solving for R	370
	21.4.4 Finding $\vec{\pi}_0$	371
	21.4.5 Performance Metrics	372
21.5	More Complex Chains	372
21.6	Readings and Further Remarks	376
21.7	Exercises	376

CONTENTS		xiii
22	Networks with Time-Sharing (PS) Servers (BCMP)	380
22.1	Review of Product-Form Networks	380
22.2	BCMP Result	380
22.2.1	Networks with FCFS Servers	381
22.2.2	Networks with PS Servers	382
22.3	M/M/1/PS	384
22.4	M/Cox/1/PS	385
22.5	Tandem Network of M/G/1/PS Servers	391
22.6	Network of PS Servers with Probabilistic Routing	393
22.7	Readings	394
22.8	Exercises	394
23	The M/G/1 Queue and the Inspection Paradox	395
23.1	The Inspection Paradox	395
23.2	The M/G/1 Queue and Its Analysis	396
23.3	Renewal-Reward Theory	399
23.4	Applying Renewal-Reward to Get Expected Excess	400
23.5	Back to the Inspection Paradox	402
23.6	Back to the M/G/1 Queue	403
23.7	Exercises	405
24	Task Assignment Policies for Server Farms	408
24.1	Task Assignment for FCFS Server Farms	410
24.2	Task Assignment for PS Server Farms	419
24.3	Optimal Server Farm Design	424
24.4	Readings and Further Follow-Up	428
24.5	Exercises	430
25	Transform Analysis	433
25.1	Definitions of Transforms and Some Examples	433
25.2	Getting Moments from Transforms: Peeling the Onion	436
25.3	Linearity of Transforms	439
25.4	Conditioning	441
25.5	Distribution of Response Time in an M/M/1	443
25.6	Combining Laplace and z-Transforms	444
25.7	More Results on Transforms	445
25.8	Readings	446
25.9	Exercises	446
26	M/G/1 Transform Analysis	450
26.1	The z-Transform of the Number in System	450
26.2	The Laplace Transform of Time in System	454
26.3	Readings	456
26.4	Exercises	456
27	Power Optimization Application	457
27.1	The Power Optimization Problem	457
27.2	Busy Period Analysis of M/G/1	459
27.3	M/G/1 with Setup Cost	462

xiv	CONTENTS	
27.4	Comparing ON/IDLE versus ON/OFF	465
27.5	Readings	467
27.6	Exercises	467
VII Smart Scheduling in the M/G/1		
28	Performance Metrics	473
28.1	Traditional Metrics	473
28.2	Commonly Used Metrics for Single Queues	474
28.3	Today's Trendy Metrics	474
28.4	Starvation/Fairness Metrics	475
28.5	Deriving Performance Metrics	476
28.6	Readings	477
29	Scheduling: Non-Preemptive, Non-Size-Based Policies	478
29.1	FCFS, LCFS, and RANDOM	478
29.2	Readings	481
29.3	Exercises	481
30	Scheduling: Preemptive, Non-Size-Based Policies	482
30.1	Processor-Sharing (PS)	482
30.1.1	Motivation behind PS	482
30.1.2	Ages of Jobs in the M/G/1/PS System	483
30.1.3	Response Time as a Function of Job Size	484
30.1.4	Intuition for PS Results	487
30.1.5	Implications of PS Results for Understanding FCFS	487
30.2	Preemptive-LCFS	488
30.3	FB Scheduling	490
30.4	Readings	495
30.5	Exercises	496
31	Scheduling: Non-Preemptive, Size-Based Policies	499
31.1	Priority Queueing	499
31.2	Non-Preemptive Priority	501
31.3	Shortest-Job-First (SJF)	504
31.4	The Problem with Non-Preemptive Policies	506
31.5	Exercises	507
32	Scheduling: Preemptive, Size-Based Policies	508
32.1	Motivation	508
32.2	Preemptive Priority Queueing	508
32.3	Preemptive-Shortest-Job-First (PSJF)	512
32.4	Transform Analysis of PSJF	514
32.5	Exercises	516
33	Scheduling: SRPT and Fairness	518
33.1	Shortest-Remaining-Processing-Time (SRPT)	518
33.2	Precise Derivation of SRPT Waiting Time*	521

Cambridge University Press

978-1-107-02750-3 - Performance Modeling and Design of Computer Systems: Queueing Theory in Action

Mor Harchol-Balter

Table of Contents

[More information](#)

CONTENTS		xv
33.3	Comparisons with Other Policies	523
33.3.1	Comparison with PSJF	523
33.3.2	SRPT versus FB	523
33.3.3	Comparison of All Scheduling Policies	524
33.4	Fairness of SRPT	525
33.5	Readings	529
 <i>Bibliography</i>		 531
<i>Index</i>		541