

STOCHASTIC MODELING AND THE THEORY OF QUEUES

Ronald W. Wolff

University of California, Berkeley



Prentice Hall, Englewood Cliffs, New Jersey 07632

Library of Congress Cataloging-in-Publication Data

WOLFF, RONALD W.,

Stochastic modeling and the theory of queues / Ronald W. Wolff.

p. cm.

Bibliography: p.

Includes index.

ISBN 0-13-846692-0

1. Stochastic processes—Mathematical models. 2. Queueing theory.

I. Title.

QA274.W65 1988 88-19624

519.2—dc19 CIP

Editorial/production supervision and

interior design: Patrice Fraccio

Manufacturing buyer: Mary Noonan



© 1989 by Prentice Hall, Inc.

A Division of Simon & Schuster

Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-846692-0

PRENTICE-HALL INTERNATIONAL (UK) LIMITED, *London*

PRENTICE-HALL OF AUSTRALIA PTY. LIMITED, *Sydney*

PRENTICE-HALL CANADA INC., *Toronto*

PRENTICE-HALL HISPANOAMERICANA, S.A., *Mexico*

PRENTICE-HALL OF INDIA PRIVATE LIMITED, *New Delhi*

PRENTICE-HALL OF JAPAN, INC., *Tokyo*

SIMON & SCHUSTER ASIA PTE. LTD., *Singapore*

EDITORA PRENTICE-HALL DO BRASIL, LTDA., *Rio de Janeiro*

PREFACE

Many years ago, I started to write a first-year graduate-level book on stochastic processes and the theory of queues, for two main reasons:

- (1) Expositions of renewal theory began with “hard stuff,” e.g., Blackwell’s Theorem. Time-average behavior (the “easy stuff”), where most of the applications are, was hardly mentioned. Even when the easy stuff was mentioned, e.g., in a problem at the end of a chapter, it was not pointed out that the easy stuff does not depend on the hard stuff.
- (2) Expositions of queueing theory were in even worse shape. The subject seemed to be a catalog of models that were solved by unrelated and often highly technical mathematical methods. Messy mathematical expressions were *derigueur*. Unifying ideas were missing.

My early treatment of the easy stuff, which was based on Tauberian theorems, wasn’t so easy, and while I had most of the right unifying ideas for queueing theory, my understanding of them wasn’t what it is today. Partially for these reasons, this first effort fizzled.

Some years later, Stuart Dreyfus and I decided to write an introductory operations research text, where he would handle optimization, and I would handle stochastic models. The treatment of stochastic models in most introductory texts is little more than formula plugging. I was confident that unifying ideas can be explained at the undergraduate level.

For this purpose, the time-average approach was simplified. A basic idea was to envisage processes evolving over time, either by drawing pictures, or

observing real processes that we were modeling. A time average then became an average over sample paths of the process. Because all the models we would encounter were regenerative, and had renewal processes embedded in them, we needed only the strong law of large numbers to derive time-average properties. With this approach, the inspection paradox and related effects became easy to understand. Tauberian theorems were banished!

This fit beautifully with fundamental queueing results such as “ $L = \lambda w$,” a relation between continuous- and discrete-time averages. The concept of work was introduced, and what is now called the PASTA property of the Poisson process was explained and applied. Elementary queueing models were compared in important and understandable ways. At that time, I wrote most of what now are the more elementary parts of Chapter 5.

A publisher who was very interested in our project sent us several reviews. The reviewers all held the view that we were writing two incompatible books. The problem seemed to be that I was after deep conceptual understanding, through a unified approach, while Dreyfus, feeling that there is no unified approach to optimization in general, was writing more of a guided tour of diverse optimization methods. Although the publisher was still enthusiastic, we agreed with the reviewers, and abandoned the joint project.

Having overcome problems that stalled my first effort, and with nearly 100 pages of new notes, I returned to my original objective. This is the result.

In the intervening years, many other books have come out, and of course we know more now. In fact, several of the topics covered in latter chapters have exploded into subfields. This growth has made the need for a unified approach even more important today than it was when my first effort began.

While written as a graduate text, this book contains a substantial amount of material not found in other general-purpose queueing books. In addition, sections at the end of seven chapters discuss important points in detail, and relate them to the literature. These features should make this book a useful general reference both for research and applications.

Most chapters have extensive sets of problems. Some of them extend the theory, while others supply it, sometimes in artful ways. Many involve problem formulation, i.e., the ability to translate into appropriate mathematics a problem stated in physical terms. This is the reason for the word *modeling* rather than either *models* or *processes* in the title.

While this has turned out to be a larger book than planned, it is my hope that the unified approach and the stress on intuitive ideas have not been diluted by its size, and that they will improve our understanding of stochastic phenomena.

We strive for rigor as well as intuitive understanding, but we also have placed a limit on assumed mathematical background that sometimes makes this difficult to achieve, e.g., in the definition of what are called stopping times. These instances occur rarely, and the benefit of this limit is greater accessibility. Furthermore, it is my firm belief that a thorough understanding of the stochastic phenomena in this book does not require a higher level of mathematical abstraction.

It is a pleasure to acknowledge the corrections and suggestions received from many friends, colleagues, students, and former students. Betsy Greenberg, George Shanthikumar, Karl Sigman, and Ward Whitt looked over the entire manuscript, and Gordon Newell recently used portions of the manuscript in a course. Their suggestions have led to numerous improvements. I also thank Philip Bitar and Yat-Wah Wan for numerous corrections and other suggestions.

Preparation of this manuscript would have been impossible without the generous technical support provided by the Department of Industrial Engineering and Operations Research at Berkeley. Finally, I thank Genji Yamazaki and the Tokyo Metropolitan Institute of Technology for their support during the editorial phase of the manuscript.

Ronald W. Wolff

ON USING THIS BOOK

As a text, this book can be used in several ways. There is a natural division: Chapters 1 through 4 for a renewal-theory (regenerative process) based course on stochastic processes, and the rest of the book for a course on queueing theory. In the latter case, there is more material than can be covered in one course. Many of the details of Chapter 11 may be omitted, as may whole topics such as Chapter 7, but it is my preference to at least explain what every major topic is about. For example, Chapter 7 helps put Chapter 6 in perspective.

At the cost of some redundancy, the introduction to queueing theory in Chapter 5 is written so that fundamental ideas can be understood without detailed knowledge of earlier chapters. Subsequent chapters rely more heavily on the earlier ones. Chapters 6 through 8 can be read independently of each other, and, aside from some knowledge of renewal theory, Chapter 9 is nearly self-contained.

The minimum prerequisites for this book are a year of calculus and a calculus-based course in probability theory. For Markov chains, some knowledge of matrix algebra is also helpful.

From time to time, we use mathematical concepts and results that go somewhat beyond these prerequisites. These include different definitions of an integral, and the validity of interchanging operations such as integral and sum. These issues arise primarily in proofs, rather than in examples or problems, and are discussed in the Appendix.

Shortcomings in probability background are more difficult to pin down. Inadequate practice at problem solving generally, and in conditional arguments in particular, is common. Thus one often hears "I thought I understood the lecture, but I can't do the homework." Certain elementary but useful matters of technique, e.g., indicator functions, may be unfamiliar. In stochastic processes, it is often very useful to fix a point in the sample space, so that a sequence of random

variables becomes a sequence of real numbers. This may also seem new, but it is important for a proper understanding of the strong law of large numbers.

Chapter 1 is intended to help remedy for these shortcomings. It is intended as a review. Topics of particular importance for the book that may also be unfamiliar are treated in some detail toward the end of this chapter. For first-year graduate students, a heterogenous group, I find that at least a week of review of this material is worthwhile.

We occasionally footnote sections, portions of sections, or even portions of proofs with an asterisk (*). The reason, often given in the text, is one or more of these: The mathematical level is higher, this is a special topic that may be omitted, or what follows is intricate and does not contribute to understanding the results.

INTRODUCTION

In this book, we strive for a balance between the theory of stochastic processes and applications of the theory to model building and the analysis of specific models. The term *stochastic modeling* is intended to convey the latter emphasis.

Much of the theory is based on the notion of *starting over* of what we call *regeneration cycles*. This is the main concept underlying renewal theory; it provides the basis for our analysis of regenerative processes in general, and discrete- and continuous-time Markov chains in particular. Nearly all the queueing models in this book are regenerative, and this approach provides a theoretical foundation for these models as well.

Because this concept is so basic, it is explicitly introduced and used early in Chapter 2, much earlier than other books of comparable level.

Renewal theory is primarily concerned with *asymptotic* or *limiting* behavior. We place great emphasis on *time averages*, which have the technical name of *Cesàro limits*. For example, we ask “what fraction of time is a process in state B ,” rather than the corresponding limiting probability of the process being in state B as time approaches infinity. We call limits of the latter type *pointwise*.

There are two reasons for this emphasis: First, a fraction of time is generally of greater practical significance than the corresponding “limit at infinity.” It has a more meaningful connection with performance measures for the models we will build. Secondly, the theory underlying the existence and representation of time averages is more elementary, and the averages hold under more general conditions than in the case for the corresponding pointwise limits.

Related to time averages is the notion of observing an ongoing process “at random,” where more formally, we consider stationary versions of processes.

The time-average point of view helps us to think in a clear and consistent manner, and to avoid mistakes. For example, the *inspection paradox* and related batch effects, which at first seem counterintuitive, are introduced early in Chapter 2 and carefully explained from the time-average point of view. The objective is to improve intuition so that when these effects occur in more complex models later (or, for that matter, in everyday life!), intuition leads in the right direction.

Realizations (sample paths) of elementary processes are shown in several figures. It is important to be able to think in these terms and to visualize processes evolving over time. This, after all, is what we would see observing the real thing—realized values of quantities that we model as random variables.

Theorems about the existence of pointwise limits, notably Blackwell's Theorem and the Key Renewal Theorem, are important for a full theoretical treatment of renewal theory. Sometimes they are relevant in applications. This topic is also treated in Chapter 2, but separately from and after time averages. Except for periodic chains in Chapter 3, this more advanced material is not needed in the rest of this book, and may be omitted. (Graduate students should be aware of these results and understand through elementary examples that pointwise limits may not exist in the lattice case.)

Another reason for the clear separation of time averages and pointwise limits is this: By mixing the two, or what is worse, hardly mentioning time averages, the reader is likely to come away with the mistaken impression that distributions must be *nonlattice* or *spread out* (see Chapter 2) for certain results to hold. This impression is reinforced by far too many papers in the literature that make these assumptions when they are not necessary. Thirty lashes!

Our point of view is the same for discrete-time processes, notably discrete-time Markov chains (Chapter 3). The distinction between periodic and aperiodic states is the analogue of lattice versus nonlattice distributions. A time average is now an average over discrete time points, usually the integers, and in Markov chains, a fraction of time becomes a *fraction of transitions*. By formulating in terms of time averages, an important theorem about the existence and uniqueness of a stationary probability vector is proven without mention of the distinction between periodic and aperiodic states.

Theoretically more difficult results about pointwise limits for discrete-time chains are covered toward the end of Chapter 3, and may be omitted.

Continuous-time Markov chains are treated in Chapter 4, along with (briefly) semi-Markov processes. New theoretical difficulties can occur for these chains, notably the possibility of an infinite number of transitions in finite time, called an *explosion*. This possibility will not occur (with positive probability) for applications of these chains in the rest of the book. Consequently, our approach is to introduce sufficient conditions that both eliminate this possibility and cover our subsequent needs. Both here and in Chapter 3, we briefly treat what are called *stationary measures*, because they are useful in Chapter 6.

Many examples in Chapters 2 through 4 and problems at the end of each chapter emphasize modeling for a wide range of applications, including queues,

inventories, transportation, replacement and maintenance policies, and reliability. The organization of these chapters, however, is based on theoretical considerations: classes of stochastic processes and their properties.

Queueing theory, which is the subject of the rest of the book, is about modeling and analysis of congestion and delays at service facilities. Nearly all of the models we discuss are stochastic processes that fall under one or more headings in earlier chapters. However, queueing theory cannot be systematically treated by a series of examples in chapters on stochastic processes. Nor should it be viewed as a “handbook” of models that, as we proceed from chapter to chapter, are treated by what may appear to be unrelated mathematical methods.

We present a unified treatment of queueing theory, based on the time-average point of view introduced in Chapter 2.

Chapter 5 is an overview, a queueing short course. Measures of performance such as L and w are defined as (continuous or discrete) time averages. Important concepts and results are introduced and applied as early as possible, including “ $L = \lambda w$,” the concept of *work-in-system* and its relation to other quantities, and the result *Poisson arrivals see time averages (PASTA)*. These results are treated in a more theoretical manner toward the end of the chapter. Queues that can be modeled as continuous-time Markov chains are included for two purposes: to show how to model complicated situations, and to facilitate the comparison of single and multichannel queues, a theme we return to several times in later chapters.

The remaining chapters are more specialized, and each has a substantial literature. We will not attempt to summarize them here; the introduction and the last section of each chapter should meet this need.

Except in Chapter 5, where we briefly consider rush-hour behavior, our queueing models are regenerative processes. See Sections 2-21 and 11-6 for discussion of recent results on the regenerative nature of queues once thought not to have this property. For these results, it is necessary to generalize how we define a regenerative process; see Section 2-21. While beyond our scope, properties of these more general regenerative processes and their application to queueing networks is currently an active research area.

CONTENTS

PREFACE	xiii
INTRODUCTION	xvii
1 PROBABILITY THEORY	1
1-1 Interpretations of Probability	1
1-2 The Probability Space and the Algebra of Events	6
1-3 Probability Measure	7
1-4 Conditional Probability; Bayes' Theorem	9
1-5 Random Variables	12
1-6 Discrete Distributions	13
1-7 Continuous Distributions	14
1-8 Joint Distributions	16
1-9 Mathematical Expectation; Magnitude and Dispersion of Random Variables	18
1-10 Sums of Random Variables; Correlation, Independence, and Convolution	23

1-11	Functions of Random Variables; Transformations	26
1-12	Generating Functions and Other Integral Transforms	28
1-13	Conditional Distributions and Conditional Expectation	31
1-14	Applications of Indicator Functions	35
1-15	Laws of Large Numbers	38
1-16	Modes of Convergence; The Central Limit Theorem	42
	Problems	44

2 RENEWAL THEORY AND REGENERATIVE PROCESSES

52

2-1	Definition of a Stochastic Process; Realizations	53
2-2	Renewal Processes	55
2-3	Reward and Cost Models; Cumulative Processes	59
2-4	Observing a Process at Random; The Equilibrium Distribution and the Inspection Paradox	65
2-5	The Discrete Case—Batch Effects	68
2-6	The Poisson Process	69
2-7	Stopping Times for a Poisson Process	77
2-8	The Connection Between Time Averages and Pointwise Limits	79
2-9	Point Processes; Generalizations of Poisson Processes	81
2-10	Defective Renewal Processes	87
2-11	Regenerative Processes	89
2-12	The Relation between a Time Average and the Mean of a Limiting Distribution	92
2-13	Wald's Equation	97

2-14	Proof of the Expected-Value Version of the Elementary Renewal Theorem	100
2-15	The Renewal Equation for Ordinary Renewal Processes; Transient Behavior	102
2-16	The Renewal Equation for General Renewal Processes; Stationary Regenerative Processes	108
2-17	A Direct Method for Finding Transient Expressions	113
2-18	Lattice Distributions; Examples Where Pointwise Limits Do Not Exist	115
2-19	Blackwell's Theorem, the Key Renewal Theorem, and Pointwise Limits	116
2-20*	Asymptotic Normality and the Regenerative Method	121
2-21	Summary and Literature Guide	125
	Problems	130

3 DISCRETE MARKOV CHAINS

148

3-1	Discrete Markov Chains; Transition Probabilities	149
3-2	Connection with Renewal Theory	151
3-3	Communication Classes and Class Properties	153
3-4	Irreducible Chains—Positive States	157
3-5	Frequencies and Relative Frequencies for Positive Irreducible Chains	160
3-6	Costs and Rewards for Positive Irreducible Chains	163
3-7	Transient Behavior	165
3-8	Periodic and Aperiodic States; Pointwise Limits	172
3-9	Branching Processes	174

3-10	Reversed and Reversible Positive Irreducible Chains	180
3-11*	Stationary Measures for Recurrent Irreducible Chains	185
	Problems	191

4 CONTINUOUS-TIME MARKOV CHAINS

201

4-1	Pure-Jump Continuous-Time Chains	201
4-2	Some Examples	206
4-3	Regular Chains	207
4-4	The Birth and Death Process	213
4-5	Time and Transition Averages for Recurrent Irreducible Chains	215
4-6	The Backward and Forward Equations	218
4-7	Uniformizable Chains	219
4-8	Semi-Markov Processes	220
4-9*	Minimal Solutions to the Backward and Forward Equations	225
	Problems	229

5 INTRODUCTION TO QUEUEING THEORY

232

5-1	Concepts, Definitions, and Notation	233
5-2	The Basic Relation: $L = \lambda w$	235
5-3	Service Facilities; The Rules of the Game	238
5-4	Deterministic Queues—Continuous Approximation	240
5-5	The $M/M/1$ Queue	245
5-6	Arrival and Time Averages; FIFO Waiting Times	250
5-7	The $M/M/1$ Queue with Queue Limit	252
5-8	Birth and Death Models	253
5-9	The $M/M/c$ Queue; Comparison with Single-Channel Queues	256

5-10	The Scope of Markovian Queues	260
5-11	The $M/G/c$ Loss System—Insensitivity	271
5-12	Solving Balance Equations—Generating Functions	273
5-13	Work and the $M/G/1$ Queue	278
5-14	The Effect of Order of Service on System Performance; Priorities	282
5-15*	On $L = \lambda w$ and Generalizations of This Result	286
5-16	On Poisson Arrivals See Time Averages (PASTA)	293
5-17	Conclusions and Related Literature	297
	Problems	300

6 REVERSIBILITY AND QUEUEING NETWORKS

308

6-1	Reversible Chains	308
6-2	Tandem Queues	313
6-3	Jackson Networks	317
6-4	Computing the Throughput of Closed Jackson Networks	327
6-5	Approximating Closed Networks; The Fixed-Population-Mean (FPM) Method	328
6-6	Open Networks with General Customer Routes	333
6-7	Networks of Quasi-Reversible Stations	336
6-8	Symmetric Queues	338
6-9	Literature Notes	340
	Problems	341

7 OVERFLOWS AND RETRIALS

345

7-1	The Two-Stage Model	346
7-2	The $GI/M/\infty$ Queue	347

7-3	The $GI/M/\ell$ Loss System	349
7-4	The Two-Stage Model with Poisson Arrivals	350
7-5	The Equivalent Random Method and Hayward's Approximation	353
7-6	The $BM/G/\infty$ Queue	358
7-7	The $BM/G/\ell$ Loss System	360
7-8	Large Two-Stage Models; Heavy Traffic	362
7-9	Retrials and the Orbit Model; The RTA Approximation	367
7-10	Monotonicity and an Upper Bound	370
7-11	The Fredericks and Reisner Approximation for the Orbit Model	372
7-12	The Quality of the Orbit Model Approximations	373
7-13	The $M/G/1$ Loss System with Infinite Orbits	376
7-14	Other Models and the Literature Problems	377 379

8 THE $M/G/1$ AND $GI/M/c$ QUEUES

381

8-1	Analysis of an Embedded Markov Chain for the $M/G/1$ Queue	381
8-2	FIFO Waiting Times for the $M/G/1$ Queue	385
8-3	Time and Transition Averages	387
8-4	The $M/G/1$ Busy Period	388
8-5	The $M/G/1$ Queue with Exceptional First Service; LIFO Waiting Times	392
8-6	The $GI/M/1$ Queue	395
8-7	The $GI/M/c$ Queue Problems	398 399

9 RANDOM WALKS AND THE $GI/G/1$ QUEUE**405**

- 9-1 The $GI/G/1$ Queue 405
- 9-2 Random Walks and Ladder Variables; Basic Terminology and Notation 406
- 9-3 Duality 409
- 9-4 A Representation of the Ladder Height Distribution 413
- 9-5 Conditions for Finite Ladder Height Moments 415
- 9-6 A Combinatorial Result 416
- 9-7 The Distribution of Ladder Epochs 417
- 9-8 Stability Conditions for the $GI/G/1$ Queue; Associated Queues 419
- 9-9 Busy Periods 420
- 9-10 Conditions for Finite Delay Moments 422
- 9-11 Preemptive LIFO and Work 424
- 9-12 Wiener-Hopf Factorization; The Delay Distribution Transform 426
- 9-13 Literature Notes 430
- Problems 431

10 WORK CONSERVATION AND PRIORITY QUEUES**436**

- 10-1 Work Conservation and Other Conservation Principles 437
- 10-2 Fixed Priority Classes 439
- 10-3 Overall Delay and Cost Considerations 444
- 10-4 Priority Rules Based on Attained Service; Tagging Jobs 448
- 10-5 The $M/G/1$ Queue with Generalized Vacations 457
- 10-6 Related Literature 464
- Problems 465

11	BOUNDS AND APPROXIMATIONS	473
11-1	<i>GI/G/1</i> Delay in Queue	474
11-2	Bounds on Expected Delay for General <i>GI/G/1</i> Queues	476
11-3	Tail Properties of Distributions	480
11-4	<i>GI/G/1</i> Bounds Based on Tail Properties of the Inter-Arrival Distribution	481
11-5	Comparison Methods with Applications to the <i>GI/G/1</i> Queue	485
11-6	The <i>GI/G/c</i> Queue—Basic Properties	493
11-7	Bounds for the <i>GI/G/c</i> Queue; Delay Moments	496
11-8	Approximations for the <i>M/G/c</i> Queue; Queue Limits and Retrials	504
11-9	Tandem Queues—Some General and Light-Traffic Results	506
11-10	<i>GI/G/1</i> and <i>GI/G/c</i> Queues in Heavy Traffic	517
11-11	Bounds on Delay Distributions for <i>GI/G/1</i> Queues	519
11-12	Further Discussion and Literature Notes	520
	Problems	522
	APPENDIX	526
A-1	Sequences and Series; Limits and Convergence	526
A-2	Integration; <i>dF</i> Notation	530
A-3	Interchanging Limit and Integral	535
A-4	Differentiating Integrals and Sums	536
	REFERENCES	539
	INDEX	549