PREFACE

In 1971 we began teaching a graduate course at MIT, "Analysis of Urban Service Systems." The course was stimulated in part by the urban turmoil of the late 1960s and early 1970s, which had turned many students toward problems of social and public concern. It was also motivated by our perception of a growing gap between courses that were primarily methodological in nature and those that were "applications-oriented," that is, dealing with real-world problems. Students appeared to be having a difficult time linking the theoretical techniques to actual urban settings.

We thus attempted to design a course that provided students with both a set of relevant analytical skills and an awareness of and sensitivity to the people-related and institutional issues that arise in practice. We felt that both purposes could best be served by focusing on logistically oriented deployment problems of certain urban service systems:

- Door-to-door pickup and delivery services (refuse collection, mail delivery)
- Emergency services (police, fire, emergency medical, emergency repair)
- Transportation services (buses, subways, jitneys, taxicabs, paratransit services)
- Certain street maintenance services (snowplowing, street sweeping, and cleaning)
- Various services provided at fixed locations (libraries, little city halls, outpatient clinics, recreation centers)

- Certain home visitation services (social worker services, meals for the elderly)

Many of these services face complicated problems related to spatial and temporal deployment of limited resources. It is this class of problems that we selected for the course.

The mathematical methods and models required to address urban deployment problems are often new and exciting. Even standard methods may require special tailoring when applied in an urban context. Many models require simultaneous consideration of both temporal and spatial relationships. *Geometrical probability* ideas are necessary to model the spatial relationships between service demands and service providers. *Queueing theory* is required to incorporate congestion caused by uncertainties in time of occurrence of service demands and length and type of service required. Reflecting the fact that many queues in a city are spatially distributed, the queueing models must often be situated in a geometrical context—with spatially distributed customers and/or servers. *Network* or *graph theory* is needed in the analysis of transportation and routing problems over the streets of a city. *Simulation*, fashioned to an urban context, may be necessary when analytical techniques fail.

To help develop knowledge of the process of implementing quantitative methods in practice, we have included a chapter on implementation. We have found that this material is most appreciated by students within the context of term projects, in which small groups of students work with a local urban service agency to help formulate, and possibly solve, an actual deployment problem.

The book is written for advanced undergraduates or beginning graduate students who have had a solid one-semester course in applied probability. Based on our 10-year teaching experience with earlier versions of this material, we expect that the book will be of interest to a diverse set of students, including those in engineering, operations research and management science, public policy, urban planning, and public administration.

A primary purpose of the book is to develop skills in formulating mathematical models from word statements of physical situations. To this purpose, we have included numerous examples and suggested exercises throughout the chapters, and we have ended each methodology chapter with an average of 20 problems ranging from elementary to difficult. Most of the problems are original to this text. Although we do not recommend that any student attempt all the problems, it is essential that for each chapter covered a representative sample be attempted. (Instructors may obtain written solutions for approximately half the problems from the authors.)

Chapter 2 reviews the probability prerequisites necessary for later chapters. The chapter covers sample spaces, events, random variables, conditional probability, expectation, well-known probability mass and density functions, and worked examples. Although much of this may be review, it is strongly recommended that the student read and study the chapter to develop

a sense of our physically oriented, model-based point of view. The "pedestrian crossing" example at the end of the chapter provides a good check on one's understanding of the concepts of the chapter.

In Chapter 3 we apply probabilistic reasoning to develop the geometrically oriented methods and results that are essential to our understanding of spatially distributed urban service systems. In the first part of the chapter we focus on functions of random variables, meaning that we must derive the (joint) probability distribution of one or more random variables that are expressed as functions of one or more other random variables whose (joint) probability distribution is known. Most of the examples worked out in the chapter pertain directly to urban service systems and may be of interest in their own right. Occasionally, when the standard techniques of probabilistic analysis are cumbersome in a geometric setting, there are available a number of special geometrically oriented tools, methods, and (sometimes) tricks that may facilitate the analysis. *Geometrical probability* is the name given to this collection. Much work in this area was done in the latter part of the nineteenth century, in such diverse areas as virology, astronomy, crystallography, and forestry, all far from the urban scene. We have selected a representative sample of geometrical probability methods in the second part of the chapter to illustrate the kind of thinking that has been applied to spatially oriented problems and to develop results useful later in the text. The third part of the chapter applies the methods of the first two parts to develop some useful "rules of thumb" that interrelate travel distances, travel times, and area geometries. Several of these relationships were tested and reinforced empirically in the early 1970s by researchers at the New York City Rand Institute. The chapter concludes with analyses of various spatially distributed stochastic processes, particularly the Poisson process, which is useful in modeling the spatial distributions of demands for many urban services.

Once the methods of geometrical probabilistic analysis are firmly grounded, we proceed in Chapter 4 to a rather extensive treatment of queueing theory as it had developed in operations research up through the 1960s and early 1970s. We cover in detail Markovian "birth and death" queues, which assume Poisson arriving customers and service times that are negative exponential. We illustrate with this simple structure how to model the complexities of multiple servers, balking, reneging, finite capacity, and so on, all presented within the context of the telephone queue associated with an urban emergency service number such as "911." We also develop key results for the single-server queue, with general service times and Poisson arrivals, with applications to an ambulance deployment problem. Also covered are priority queues, simple bounds and approximations, and a brief introduction to queueing networks. While motivated from an urban perspective, we feel that the chapter is a useful, self-contained, transform-free introduction to queues.

In Chapter 5 we merge the ideas of Chapters 3 and 4 to discuss spatially distributed queues. Since there are far too many spatially distributed queues
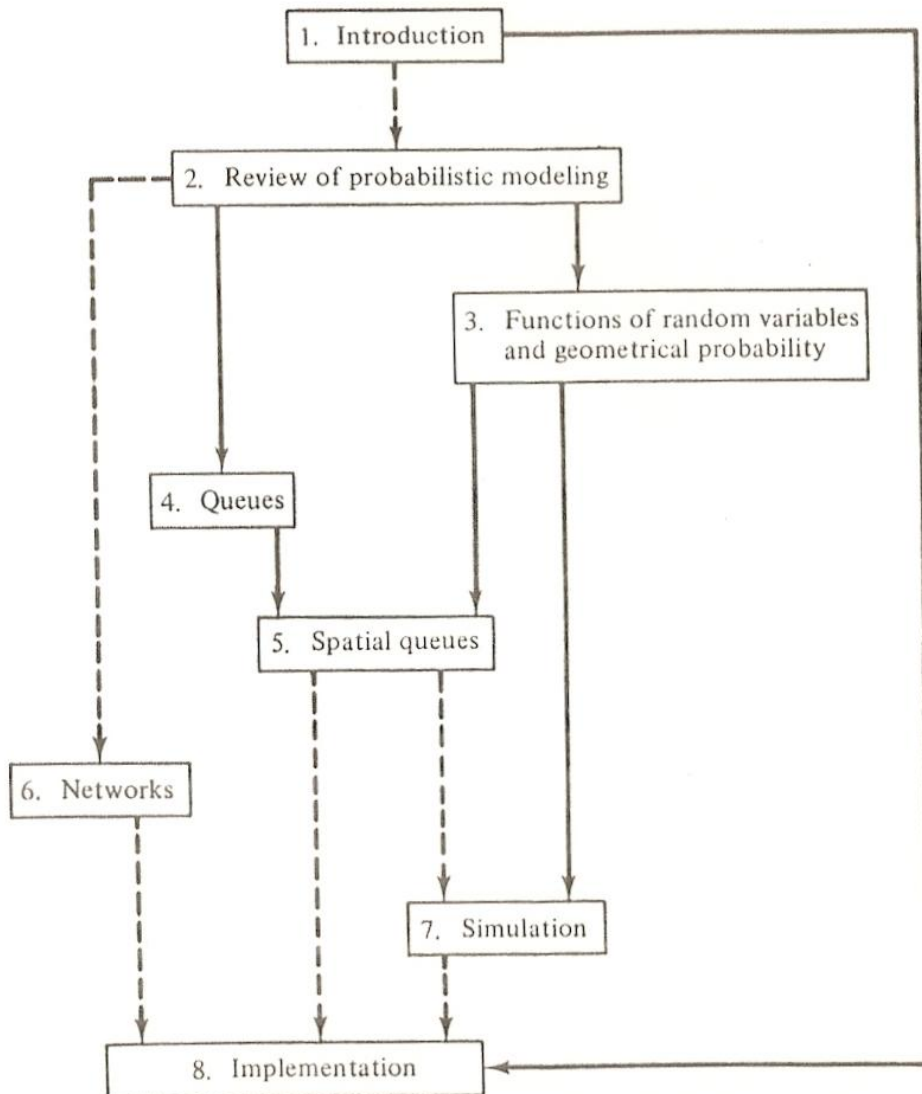
in a city to cover carefully within one chapter, we focus most of our attention on one representative class—that which models accurately urban emergency services. We develop analytical results for single-server and two-server models, but we must resort to a computer-implemented model for three or more servers. This model, called the "hypercube model" because of its state space, illustrates how a particular problem structure can be exploited for efficient computer implementation and solution. Numerous potential (and actual) applications of the hypercube model are described. When the number of spatially distributed servers becomes large, one can use a "hypercube approximation procedure," which reduces the number of required simultaneous equations for N servers from 2^N to N. This procedure is derived in the chapter and illustrated with an $N = 3$ example. The chapter concludes with some useful results for many-server spatially distributed queues.

Chapter 6 changes focus initially from probabilistic to deterministic analysis, dealing with deployment problems that occur on networks. A network is a set of nodes (e.g., street intersections, major neighborhoods, towns) connected by a set of links (e.g., street segments, major arteries, highways). Deployment problems on a network can be considered to be routing problems, or location problems, or network design problems. Many of these problems are of interest in themselves, while others provide inputs to other models. An important routing problem is the determination of a minimal-travel-time path from any node to any other node. This "minimal-travel-time-path" problem is of independent interest, and its solution can be used to provide an input travel time matrix for the hypercube model. Other routing problems include "node-covering" problems such as the famous traveling salesman problem, and "link-covering" problems such as the Chinese postman problem. Location problems deal with the siting of a finite number of facilities on a network in order to optimize some objective function. Examples include the now well-known N-median and N-center problems. A network design problem requires optimal structure of nodes and links in order to best achieve some objective; an example is the minimal-spanning-tree problem. All the examples cited are covered in Chapter 6. Moreover, the chapter merges, in a new and (we think) useful way, traditional network problems with new "vehicle-routing" problems. It also refers briefly to active research in probabilistic networks. We feel that it is a self-contained treatment of network theory, as it applies to deployment problems, and should be of interest independently of the urban context within which it is presented.

Chapter 7 discusses what one should do when and if analytical techniques fail—*simulate*. The chapter is a self-contained introduction to Monte Carlo simulation, with special emphasis given to geometrical considerations that arise when simulating spatially distributed systems. Included, for instance, are methods for generating random sample points from within a polygon of arbitrary shape and for determining whether various geometrical figures overlap.

To provide a context for the student projects that we mentioned earlier, Chapter 8 concludes the book with a discussion of *implementation*. We cannot overemphasize the importance we place on this subject. The techniques of this book are exciting and often new; they can provide many insights to urban decision makers. However, before eagerly launching out into the real world armed with these tools, the student must have an appreciation for the broader setting of any study that would utilize the methods of this book in an operating urban setting. The chapter provides thumbnail sketches of some of our own experiences, as well as those of our students. We then attempt to build from the cases to discuss implementation issues that are related both to models and to people and their institutions.

The material covered in Chapters 1–8 could easily require a two-semester course, especially if augmented by a substantial real-world project. In designing a one-semester course, the instructor should keep in mind the following precedent relationships among chapters:

```
                    ┌──────────────────────┐
                    │ 1.  Introduction      │───────────────────────┐
                    └──────────┬───────────┘                        │
                               ┆                                     │
                               ▼                                     │
        ┌ ─ ─ ─ ─ ─┌──────────────────────────────────┐             │
        ┆          │ 2.  Review of probabilistic modeling│            │
        ┆          └──────┬───────────────────┬────────┘             │
        ┆                 │                   │                      │
        ┆                 │                   ▼                      │
        ┆                 │       ┌────────────────────────────┐     │
        ┆                 │       │ 3.  Functions of random variables│ │
        ┆                 │       │     and geometrical probability │ │
        ┆                 │       └──────┬──────────┬────────────┘   │
        ┆                 ▼              │          │                │
        ┆          ┌───────────┐         │          │                │
        ┆          │ 4.  Queues │         │          │                │
        ┆          └─────┬─────┘          │          │                │
        ┆                │                │          │                │
        ┆                ▼                │          │                │
        ┆          ┌────────────────┐     │          │                │
        ┆          │ 5.  Spatial queues│   │          │                │
        ┆          └──┬────┬────┬────┘     │          │                │
        ▼             ┆    ┆    ┆          │          │                │
   ┌──────────────┐   ┆    ┆    ┆          │          │                │
   │ 6.  Networks  │   ┆    ┆    ┆          ▼          │                │
   └──────┬───────┘   ┆    ┆    ┆   ┌───────────────┐  │                │
          ┆           ┆    ┆    ┆   │ 7.  Simulation │  │                │
          ┆           ┆    ┆    ┆   └───────┬───────┘  │                │
          ▼           ▼    ▼    ▼           ▼          │                │
   ┌────────────────────────────────────────────────┐ │                │
   │ 8.  Implementation                              │◄─────────────────┘
   └────────────────────────────────────────────────┘
```

(A dashed line indicates weaker dependence than a solid line)

In our MIT one-semester offering of the subject "Analysis of Urban Systems" (now renamed "Urban Operations Research"), we have found it possible to cover the key ideas in each of the eight chapters. In any one-semester offering, however, advanced material in Chapters 3–6 will have to be selectively chosen, with some of the material left for student reading (e.g., the hypercube approximation procedure in Chapter 5, queueing bounds and approximations and priority queues in Chapter 4, Crofton's method in Chapter 3, the p-center problem in Chapter 6). By selecting Chapter 2, functions of random variables in Chapter 3, and Chapters 4, 6, and 7, the instructor can cover all the introductory methods of operations research except linear programming.

<div style="text-align: right">

Richard C. Larson
Amedeo R. Odoni

</div>

Cambridge, Massachusetts