Chapter 1 : Optimal Design of Queueing Systems - CRC Press Book

The First Comprehensive Book on the Subject. Focusing on the underlying structure of a system, Optimal Design of Queueing Systems explores how to set the parameters of a queueing system, such as arrival and service rates, before putting it into operation.

The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint. Except as permitted under U. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers. For permission to photocopy or use material electronically from this work, please access www. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged. Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe. Stability and Chaos 83 3. Optimal Solutions 4. Dynamic Algorithms 4. Class Dominance 4. Parallel Queues 6. Optimal Solutions 8. Dynamic Adaptive Algorithms 8. Homogeneous Waiting Costs 8. Preemptive-Resume Scheduling Rules A. Uniform Value Distribution 69 2. Work in System Preface What began a long time ago as a comprehensive book on optimization of queueing systems has evolved into two books: The distinction between design and control, admittedly, can be somewhat artificial. But the available material had outgrown the confines of a single book and I decided that this was as good a way as any of making a division. Why look at design models? In principle, of course, one can always do better by allowing the values of the decision variables to depend on the state of the system, but in practice this is frequently an unattainable goal. But such information is rarely available to a centralized controller in time to make decisions that are useful for the network as a whole. Even if it were available, the combinatorial complexity of the decision problem makes it impossible to solve even approximately in the time available. The essential difficulty with such systems is that the time scale on which the system state is evolving is comparable to, or shorter than, the time scale on which information can be obtained and calculations of optimal policies can be made. For these and other reasons, those in the business of analyzing, designing, and operating communication networks have turned their attention more and more to flow control, in which quantities such as arrival e. Indeed, many of the models, techniques, and results discussed in this book were inspired by research on flow and routing control that has been reported in the literature on communication networks. Of course, flow control is still control in the sense that decision variables can ix x PREFACE change their values in response to changes in the state of the system, but the states in question are typically at a higher level, involving congestion averages taken over time scales that are much longer than the time scale on which such congestion measures as queue lengths and waiting times are evolving at individual service facilities. For this reason, I believe that flow control belongs under the broad heading of design of queueing systems. I have chosen to frame the issues in the general setting of a queueing system, rather than specific applications such as communication networks, vehicular traffic flow, supply chains, etc. I believe strongly that this is the most appro- priate and effective way to produce applicable research. It is a belief that is consistent with the philosophy of the founders of operations research, who had the foresight to see that it is the underlying structure of a system, not the physical manifestation of that structure, that is important when it comes to building and applying mathematical models. Unfortunately, recent trends have run counter to this philosophy, as more and more research is done within a particular application discipline and is published in the journals of that discipline, using the jargon of that discipline. The result has been compartmentalization of useful research. Important re- sults are sometimes rediscovered in, say, the communication and computer science communities, which have been well known for decades in, say, the traffic-flow community. I blame the research funding agencies, in part, for this trend. These proposals and papers may begin with a detailed description of a particular application in which congestion occurs, in order to establish the credibility of the authors within the appropriate research community. Another of my basic philosophies has been to present the various models in a unified notation and terminology and, as much as possible, in a unified ana- lytical framework. In keeping with my belief expressed above that queueing theory, rather than any one or several of its applications, provides the appro- priate modeling basis for this field, it is natural that I should have adopted the notation and terminology of queueing theory. Providing a unified ana- lytical framework was a more difficult task. In the literature optimal design problems for queueing systems have been solved by a wide variety of analyt- ical techniques, including classical calculus, nonlinear programming, discrete optimization, and sample-path analysis. My desire for unity, together with space constraints, led me to restrict my attention to problems that can be solved for the most part by classical calculus, with some ventures into elemen- tary nonlinear programming to deal with constraints on the design variables. A side benefit of this self-imposed limitation has been that, although the book PREFACE xi is mathematically rigorous I have not shied away from stating results as the- orems and giving complete proofs , it should be accessible to anyone with a good undergraduate education in mathematics who is also familiar with el- ementary queueing theory. The downside is that I have had to omit several interesting areas of queueing design, such as those involving discrete decision variables e. I plan to include many of these topics in my queueing control book, however, since they are relevant also in that context. The emphasis in the book is primarily on qualitative rather than quanti- tative insights. A recurring theme is the comparison between optimal designs resulting from different objectives. An example is the by-now-classical result that the individually optimal arrival rate is typically larger than the socially optimal arrival rate. As a general principle, this concept is well known in welfare economics. Indeed, a major theme of the research on queueing design has been to bring into the language of queueing theory some of the important issues and qualitative results from economics and game theory the Nash equilibrium being another example. As a consequence this book may seem to many readers more like an economics treatise than an operations research text. I have always felt that students and practitioners would benefit from an infusion of basic eco- nomic theory in their education in operations research, especially in queueing theory. Much of the research reported in this book originated in vehicular traffic- flow theory and some of it pre-dates the introduction of optimization into queueing theory in the s. As such, it may be construed as a subtopic in nonlinear pro- gramming. An emphasis in this branch of traffic-flow theory has been on com- putational techniques and results. Although models for optimal design of queueing systems using my broad definition have proliferated in the four decades since the field began, I was surprised at how often I found myself developing new results because I could not find what I wanted in the literature. Perhaps I did not look hard enough. The proliferation of research on queueing design, together with the explosion of different application ar- eas each with its own research community, professional societies, meetings, and journals, have made it very difficult to keep abreast of all the important research. I have tried but I may not have completely succeeded. A word about the organization of the book: References for the models and results on optimal design of queues are usually given in an endnote the final section of the chapter , along with pointers to material not covered in the book. Acknowledgements I would like to thank my editors at Chapman Hall and CRC Press in London for their support and patience over the years that it took me to write this book. I particularly want to thank Fred Hillier for introducing me to the field of optimization of queueing systems a little over forty years ago. I am grateful to my colleagues at the following institutions where I taught courses or gave seminars covering the material in this book: Finally, my wife Carolyn deserves special thanks for finding just the right combination of encourage- ment, patience, and at appropriate moments prodding to help me bring this project to a conclusion. What distinguishes an optimal design model from a traditional descriptive model is the fact that some of the parameters are subject to decision and that this decision is made with explicit attention to economic considerations, with the preferences of the decision maker s as a guiding principle. The basic distinctive components of a design model are thus: Decision variables may include, for example, the arrival rates, the service rates, and the queue disciplines at the various service facilities. Typical benefits and costs include rewards to the customers from being served, waiting costs incurred by the customers while waiting for service, and costs to the facilities for providing the service. These benefits and costs may be brought together in an objective function, which quantifies the implicit trade-offs. For example, increasing the service rate will result in less time spent by the customers waiting and thus a lower waiting cost , but a higher service cost. The nature of the objective function also depends on the horizon finite or infinite , the presence or absence of discounting, and the identity of the decision maker e. Our goal in this chapter is to provide a quick introduction to these ba- sic components of a design model. We shall illustrate the effects of different reward and cost structures, the trade-offs captured by different objective func- tions, and the effects of combining different decision variables in one model. This will allow us to do the optimization with the simple and familiar tools of differential calculus. Later chapters will elaborate on each of the models introduced in this chapter, relaxing distributional assumptions and considering more general cost and reward structures and objective functions. These more general models will require more sophisticated analytical tools, including linear and nonlinear programming and game theory. We begin this chapter Sections 1. In the first example the decision variable is the service rate and in the second, the arrival rate. The simple probabilistic and cost structure makes it possible to use classical calculus to derive analytical expressions for the optimal values of the design variables. The next three sections consider problems in which more than one design parameter is a decision variable. Here a simple analysis based on calculus breaks down, since the objective function is not jointly concave and therefore the first-order optimality conditions do not identify the optimal solution. This will be a recurring theme in our study of optimal design models, and we shall explore it at length in later chapters. Again the objective function is not jointly concave and the first-order optimality conditions do not identify the optimal arrival rates. Indeed, the only interior solution to the first-order conditions is a saddle-point of the objective function and is strictly dominated by both boundary solutions, in which only one class has a positive arrival rate. Finally, in Section 1. A final word before we start. In a design problem, the values of the decision variables, once chosen, cannot vary with time nor in response to changes in the state of the system e. Design problems have also been called static control problems, in contrast to dynamic control problems in which the decision variables can assume different values at different times, depending on the observed state of the system. In the literature a static control problem is sometimes called an open-loop control problem, whereas a dynamic control problem is called a closed-loop control problem. We shall simply use the term design for the former and control for the latter type of problem. Service times are independent of the arrival process and i. A machine center in a factory: Performance Measures and Trade-offs. Typical performance measures are the number of customers in the system or in the queue and the waiting time of a customer in the system or in the queue. If the system operates for a long time, then we might be interested in the long-run average or the expected steady-state number in the system, waiting time, and so forth. All these are measures of the level of congestion. But, in all real systems, increasing the service rate costs something. One way to capture this trade-off is to consider a simple model with linear costs.

Chapter 2 : Optimal Design of Queueing Systems - [PDF Document]

Focusing on the underlying structure of a system, Optimal Design of Queueing Systems explores how to set the parameters of a queueing system, such as arrival and service rates, before putting it into operation. It considers various objectives, comparing individually optimal (Nash equilibrium), socially optimal, class optimal, and facility optimal flow allocations.

Chapter 3 : Optimal Design of Queueing Systems - Ebook Library Download

Optimal Design of Queueing Systems. Book Title:Optimal Design of Queueing Systems. The First Comprehensive Book on the Subject Focusing on the underlying structure of a system, Optimal Design of Queueing Systems explores how to set the parameters of a queueing system, such as arrival and service rates, before putting it into operation.

Chapter 4 : Optimal Design of Queueing Systems: 1st Edition (e-Book) - Routledge

DOWNLOAD PDF OPTIMAL DESIGN OF QUEUEING SYSTEMS

Shaler Stidham's book,optimal design of queueing systems, is an instant classic. He starts by clearly describing how to find the optimal service rate for an M/M/1 queue for a given arrival rate. I'm looking forward to the next book in the series, control of queueing systems.

Chapter 5 : Optimal Design of Queueing Systems - Shaler Stidham Jr - Bok () | Bokus

There has been limited previous work on optimal design of measurement times for stochastic systems, and particularly for the M/M/1 queue we have studied here. We have adapted a method to evaluate the autocovariance function of the queue quickly, and used this to numerically calculate the optimal designs for this queue.

Chapter 6 : Queueing theory - Wikipedia

Optimal Design of Queueing Systems 1st Edition by Shaler Stidham Jr. and Publisher Chapman and Hall/CRC. Save up to 80% by choosing the eTextbook option for ISBN: , X. The print version of this textbook is ISBN: ,

Chapter 7 : Optimal Design of Queueing Systems eBook: Shaler Stidham Jr.: calendrierdelascience.com: K

Optimal Design and Control of Finite-Population Queueing Systems Chao Deng A dissertation submitted to the faculty of the University of North Carolina at.