Preface

I don't have any solution, but I certainly admire the problem.
— Ashleigh Brilliant

I take the view, and always have done, that if you cannot say what you have to say in twenty minutes, you should go away and write a book about it.
— Kord Brabazon

Everything that happens once will never happen again. But everything that happens twice will surely happen a third time.
— Paulo Coelho

From A Priori to Online Optimization Optimization systems traditionally have focused on a priori planning and are rarely robust to disruptions. The airline industry, a sophisticated user of optimization technology, solves complex fleet assignment, crew scheduling, and gate allocation problems as part of its operations using some of the most advanced optimization algorithms available. Yet unexpected events such as weather conditions or strikes produce major disruptions in its operations. In August 2005, British Airways took three to four days to resume its normal operations after a strike of one of its catering subcontractors; many of its planes and crews were at the wrong places because of the airline's inability to anticipate and quickly recover from the event. The steel industry, also a significant customer of optimization technology, typically separates strategic planning (which orders to accept) and tactical planning (which priorities to assign them) from daily scheduling decisions. The strategic decisions are based on coarse approximations of the factory capabilities and forecasts, sometimes leading to infeasible schedules, missed deadlines, and poor decisions when novel incoming orders arrive online.

Fortunately, the last decades have witnessed significant progress in optimization and information technology. The progress in speed and functionalities of optimization software has been simply amazing. Advances in telecommunications, such as the global positioning system (GPS), sensor and mobile networks, and radio frequency identification (RFID) tags, enable organizations to collect a wealth of data on their operations in real time.

It also is becoming increasingly clear that there are significant opportunities for optimization algorithms that make optimization decisions online. Companies such as UPS have their own meteorologists and political analysts to adapt their operations and schedules online. Pharmaceutical companies must schedule drug design projects with uncertainty on success, duration, and new developments. Companies such as Wal-Mart now try to integrate their supply chains with those of theirs suppliers, merging their logistic systems and replenishing inventories dynamically.

As a consequence, we may envision a new era in which optimization systems will not only allocate resources optimally: they will react and adapt to external events effectively under time constraints, anticipating the future and learning from the past to produce more robust and effective solutions. These systems may deal simultaneously with planning, scheduling, and control, complementing a priori optimization with integrated online decision making.

Online Stochastic Combinatorial Optimization This book explores some of this vision, trying to understand its benefits and challenges and to develop new models, algorithms, and applications. It studies *online stochastic combinatorial optimization* (OSCO), a class of optimization applications where the uncertainty does not depend on the decision-making process. OSCO problems are ubiquitous in our society and arise in networking, manufacturing, transportation, distribution, and reservation systems. For instance, in courier service or air-taxi applications, customers make requests at various times and the decision-making process must determine which requests to serve and how under severe time constraints and limited resources.

Different communities approach new classes of problems in various ways. A problem-driven community studies individual applications and designs dedicated solutions for each of them. A theoretically oriented community often simplifies the applications to identify core algorithmic problems that hopefully are amenable to mathematical analysis and efficient solutions. These approaches are orthogonal and often produce invaluable insights into the nature of the problems. However, many professionals in optimization like to say that "there are too many applications with too many idiosyncratic constraints" and that "an approximated solution to a real problem is often preferable to an optimal solution to an approximated problem." As a result, this book takes a third, engineering-oriented, approach. It presents the design of abstract models and generic algorithms that are applicable to many applications, captures the intricacies of practical applications, and leverages existing results in deterministic optimization.

Online Anticipatory Algorithms More precisely, to tackle OSCO applications, this book proposes the class of *online anticipatory algorithms* that combine online algorithms (from computer science) and stochastic programming (from operations research). Online anticipatory algorithms assume the availability of a distribution of future events or an approximation thereof. They take decisions during operations by solving deterministic optimization problems that represent possible realizations of the future. By exploiting insights into the problem structure, online anticipatory algorithms address the time-critical nature of decisions, which allows for only a few optimizations at decision time or between decisions.

The main purpose of this book is thus to present online anticipatory algorithms and to demonstrate their benefits on a variety of applications including online packet scheduling, reservation systems, vehicle dispatching, and vehicle routing. On each of these applications, online anticipatory algorithms are shown to improve customer service or reduce costs significantly compared to oblivious algorithms that ignore the future. The applications are diverse. For some of them, the underlying optimization problem is solvable in polynomial time. For others, even finding optimal solutions to the deterministic optimization where all the uncertainty is revealed is beyond the scope of current optimization software. Moreover, these applications capture different types of decisions. On some of them, the issue is to choose which request to serve, while, on others, the question is how to serve the request. On some of these applications, it is not clear even what the decisions should be in an online setting, highlighting some interesting modeling issues raised by online applications. In particular, the book presents some results on vehicle-routing strategies that were amazingly counterintuitive at first and seem natural retrospectively.

Of course, in practice, not all applications come with a predictive model of the future. The book also studies applications in which only the structure of model or historical data is available. It shows how to integrate machine learning and historical sampling into online anticipatory algorithms to address this difficulty.

From Practice to Theory and Back Demonstrating the benefits of online anticipatory algorithms on a number of applications, however diverse and significant, is hardly satisfying. It would be desirable to identify the class of applications that are amenable to effective solutions by online anticipatory algorithms. Such characterizations are inherently difficult, however, even in deterministic optimization: indeed optimization experts sometimes disagree about the best way to approach a novel application. OSCO applications further exacerbate the issue by encompassing online and stochastic elements.

This book attempts to provide some insights about the behavior of online anticipatory algorithms by identifying assumptions under which they deliver near-optimal solutions with a polynomial number of optimizations. At the core of online anticipatory algorithms lies an anticipatory relaxation that removes the interleaving of decisions and observations. When the anticipatory relaxation is tight, a property called *ϵ-anticipativity*, online anticipatory algorithms closely approximate optimal, a posteriori solutions. Although this is a strong requirement, the applications studied in this book are shown to be ϵ-anticipative experimentally. The inherent temporal structure of these applications, together with well-behaved distributions, seems to account for this desirable behavior. The analysis presented here is only a small first step and much more research is necessary to comprehend the nature of OSCO applications and to design more advanced online anticipatory algorithms.

A Model for Sequential Decision Making The theoretical analysis has an interesting side effect: it highlights the common abstract structure that was buried under the idiosyncracies of the applications, their models, and their algorithms. It also establishes clear links with *Markov Decision Processes* (MDP), a fundamental approach to sequential decision making extensively studied in artificial intelligence and operations research. Like MDPs, online stochastic combinatorial optimization alternates between decisions and observations, but with a subtle difference: the uncertainty in MDPs is endogenous and depends on the decider's actions. It is thus possible to define a variant of MDPs, called *Markov Chance-Decision Processes* (MCDPs), that captures online stochastic combinatorial optimization and whose uncertainty is exogenous. In MCDPs, the decision process alternates between observing uncertain inputs and deterministic actions. In contrast, the decision process in traditional MDPs, called *Markov Decision-Chance Processes* (MDCPs) here, alternates between actions whose effects are uncertain and observations of the action outcomes. As a consequence, MCDPs crystallize the essence of online stochastic combinatorial optimization that can be summarized as follows:

Anticipatory Relaxation: Contrary to MDCPs, MCDPs naturally encompass an anticipatory relaxation for estimating the future. The anticipatory relaxation can be approximated by the solutions of deterministic optimization problems representing scenarios of the future.

Online Anticipatory Algorithms: MCDPs naturally lead to a class of online anticipatory algorithms taking decisions online at each time step using the observations and approximations to the anticipatory relaxations.

Anticipativity: When the anticipatory relaxation is ϵ-anticipative, online anticipatory algorithms produce near-optimal solutions with a polynomial number of optimizations.

Learning: The distribution of the inputs can be sampled and learned independently from the underlying decision process, providing a clean separation of concerns and computational benefits.

So, perhaps now that this book is written, twenty minutes are sufficient to describe its contents, making its writing yet another humbling experience. However, the open issues it raises are daunting both in theory and practice.

Acknowledgments As Hervé Gallaire (and some other famous people) like to say, "we got by with a little help from our friends." Actually, more than a little help. First, we would like to express our gratitude to Eli Upfal and Luc Mercier who coauthored two chapters of this book. Eli is chair of the Department of Computer Science at Brown University and still finds time to talk to us. Luc Mercier has the annoying habit of asking all the difficult questions and the good habit of coming up with most of the answers. Both of them have improved our understanding of this topic tremendously. The research leading to this book started in the context of a National Science Foundation (NSF) ITR Award (DMI-0121495) in collaboration with Georgia Tech., MIT, and SUNY Buffalo. It is a pleasure to thank Shabbir Ahmed, Ismael De Farias, Michel Goemans, and George Nemhauser for all their insights and enthusiasm for research, and, of course, NSF for making it all possible. Several students and colleagues also contributed significantly to this research. Aris Anagnostopoulos is a master in debugging theoretical results and is the "man" in case you ever need — or want — to travel to Acapulco. Yannis Vergados started the research on online multiknapsacks and helped us finish it, although his passion is sport scheduling. Irit Katriel came up with the beautiful (but negative) result on suboptimality approximations. Laurent Michel has been a patient sounding board for the last five years. As always, many thanks to Bob Prior, our MIT Press editor, who has become an expert in the art of tightening and relaxing deadlines under uncertainty. Mel Goldsipe and her editors have been amazingly helpful, even correcting formulas! And, last but not least, to our families without whom this book would have remained only a possible realization of an uncertain future.