

Preface

Why this book now?

A key direction for research in systems and control involves *engineering systems*. These are highly distributed collective systems with humans in the loop. *Highly distributed* means that decisions, information, and objectives are distributed throughout the system. *Humans in the loop* implies that the *players* have bounded rationality and limited computation capabilities. In addition, decisions may also be influenced by societal and cultural habits. Engineering systems emphasize the potential of control and games beyond traditional applications.

The reason why I chose to write this book now is that, within the realm of engineering systems, a key point is the use of *game theory* to design incentives to obtain *socially desirable behaviors* on the part of the players. As an example, in *demand side management*, an increase of the electricity price on the part of the network operator may induce a change in the consumption patterns on the part the prosumers (producers-consumers). In *opinion dynamics*, sophisticated marketing campaigns may influence the market share, assuming that the customers are susceptible players sharing opinions with their neighbors. In *pedestrian flow*, informing the pedestrians on the congestion at different locations may lead to a better redistribution of the traffic. These are only some of the applications discussed in this book.

In this context, *game theory* offers a rich set of model elements, solution concepts, and evolutionary notions. The model elements are the players, the action sets, and the payoffs; the solution concepts include the Nash equilibrium, the Stackelberg equilibrium, and Pareto and social optimality; evolutionary notions shed light on the fact that equilibria are relevant only if the players can converge to such solutions in a dynamic setting. Evolutionary notions essentially turn the game into a kind of dynamic feedback system.

However, *a game theory model is more than just a dynamic feedback system*, as each player learns the environment, which in turn learns the player, and so forth. Such a coupled learning introduces a higher level of difficulty to the feedback structure.

A large portion of this book is dedicated to *games with a large number of players*. Here each player uses an aggregate description of the environment based on a distribution function on actions or states, which is the main idea in a *mean-field game*. Thus, in most examples the game is a mapping from distributions (congestion levels) to payoffs (think of the replicator dynamics).

If a game is a mapping from congestion levels to payoffs, the evolution model is a dynamic model that operates in the opposite direction: it maps flows of payoffs to flows of congestion levels. Here, *systems and control theory* provides a set of sophisticated stabilizability tools to design self-organizing and resilient systems characterized by cooperation

and competition. This book will mainly use the Lyapunov approach both in a deterministic and a stochastic setting.

Goal of this book

This book's goal is to bring together game theory and systems and control theory in the unconventional framework of engineering systems. The goal of Part I is to cover the foundations of the theory of noncooperative and cooperative games, both static and dynamic. Part I also highlights new trends in cooperative differential games, learning, approachability (games with vector payoffs), and mean-field games (large number of homogeneous players). The treatment emphasizes theoretical foundations, mathematical tools, modeling, and equilibrium notions in different environments.

The goal of Part II is to illustrate stylized models of engineered and societal situations. These models aim at providing fundamental insights on several aspects, including the individuals' strategic behaviors, scalability and stability of the collective behavior, and the influence of heterogeneity and local interactions. Other relevant issues discussed throughout the book are uncertainty and model misspecification. Remarkably, the framework of robust mean-field games is developed with an eye to grand engineering challenges such as *resilience* and *big data*.

What this book is not

This book is not an encyclopedia of game theory, and the material covered reflects my personal taste. More importantly, this book is not a collection of takeaway models and solutions to specific applications. These models need not be interpreted literally but are guidelines towards a better understanding and an efficient design of collective systems.

Structure of this book

This book is organized in two parts. Part I follows [24] and goes from Chapter 1 to 12. Chapters 1 to 4 review the foundations of noncooperative games. Chapters 5 to 6 deal with cooperative games. Chapter 7 surveys evolutionary games. Chapter 8 analyzes the replicator dynamics and provides a brief overview of learning in games. Chapter 9 deals with differential games. Chapter 10 discusses stochastic games. Chapter 11 pinpoints basics and trends in games with vector payoffs, such as *approachability* and *attainability*. Chapter 12 provides an overview of mean-field games.

Part II builds upon articles of the author and goes from Chapters 13 to 21. In particular, under the umbrella of power systems, Chapters 14 to 15 analyze demand side management and synchronization of power generators, respectively. Within the realm of sociophysical systems, Chapter 13 discusses consensus in multi-agent systems, and Chapters 16 to 18 illustrate, in order, opinion dynamics, bargaining, and pedestrian flow applications. Within the context of production/distribution systems, Chapters 19 to 21 deal with supply chain, population of producers, and cyber-physical systems.

At the end of each chapter a section entitled "Notes and references" acknowledges the work on which the chapter is based and related works.

Audience

The primary audience is students, practitioners, and researchers in different areas of Engineering such as Industrial, Aeronautical, Manufacturing, Civil, Mechanical, and Elec-

trical Engineering. However, the topic also interests scientists in Computer Science, Economics, Physics, and Biology. Young researchers may benefit from reading Part II. The comprehensive reference list enables further research. The book is self-contained and makes the path from undergraduate students to young researchers short.

Using this book in courses

This book can be used as a textbook, especially Part I. This part covers material that can be taught in first-year graduate courses. I use a tutorial style to illustrate the major points so that the reader can quickly grasp the basics of each concept.

Part I assembles the material of three graduate courses given at the Department of Mathematics of the University of Trento, at the Department of Engineering Science of the University of Oxford, and at the Department of Electrical and Electronic Engineering of Imperial College in 2013. The material has also been used for the short course given at the Bertinoro International Spring School 2015 held in Bertinoro, Forlì, Italy.

The book can also be used for an undergraduate course. To this purpose, the book is complemented with Appendix sections on mathematical review, optimization, Lyapunov stability, basics of probability theory, and stochastic stability theory. Part II shows a number of simulation algorithms and numerical examples that may help improve the coding skills of the students. The software used for the simulations is MATLAB®. *Prior knowledge* includes the material discussed in the Appendix sections.

Acknowledgments

A large part of this book is based on my research over the past 10 years. I was honored to have a number of brilliant co-authors, and I would like to mention those with whom I have worked extensively. The collaboration with Tamer Başar and Hamidou Tembine is the origin of many ideas in robust mean-field games. The collaboration with Ehud Lehrer, Eilon Solan, and Xavier Venel has inspired research on attainability (cf. Chapter 11). The joint work with Franco Blanchini is the source of several ideas on robust stabilizability of network flows appearing throughout the book. Raffaele Pesenti and Laura Giarré have helped me develop the ideas discussed in the multi-agent consensus application in Chapter 13. The bargaining model in Chapter 17 has been developed in a joint work with Angelia Nedić. The supply-chain model in Chapter 18 has been studied in a collaboration with Judith Timmer. The collaboration with Fabio Bagagiolo has inspired the design of objective functions in differential and mean-field games.

More recently, the approximation technique based on state space extension to compute mean-field equilibria has resulted from the fruitful interactions with Alessandro Astolfi and Thulasi Mylvaganam during my sabbatical at Imperial College London in 2013. The collaboration with Antonis Papachristodoulou and Xuan Zhang has inspired the pedestrian flow model in Chapter 18. My special thanks to Xuan, who has contributed the simulations in Chapter 18. I really enjoyed sharing thoughts with Mark Cannon, and the resulting ideas combining games and receding horizon are discussed in Chapter 16 in the context of opinion dynamics. The collaborations with Antonis, Xuan, and Mark started during my sabbatical period in Oxford in 2013.

Many thanks are due to the several PhD students, postdocs, and fellows who have attended the courses and have contributed to the improvement of the material with their comments and questions.

I deeply thank my sister Helga, who has been generous with love and encouragement despite the distance between us. Finally, I would like to thank Claudia for her enormous support and for sharing the ups and downs with me.

I hope you will enjoy reading the book as much as I did writing it!