



Algorithmic Game Theory

Edited by **Noam Nisan, Tim Roughgarden,
Éva Tardos, and Vijay V. Vazirani**

Foreword by **Christos H. Papadimitriou**

CAMBRIDGE

0225
A396

Algorithmic Game Theory

Edited by

Noam Nisan

Hebrew University of Jerusalem

Tim Roughgarden

Stanford University

Éva Tardos

Cornell University

Vijay V. Vazirani

Georgia Institute of Technology



E2008000518



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

Cambridge University Press

32 Avenue of the Americas, New York, NY 10013-2473, USA

www.cambridge.org

Information on this title: www.cambridge.org/9780521872829

© Noam Nisan, Tim Roughgarden, Éva Tardos, Vijay V. Vazirani 2007

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2007

Printed in the United States of America

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Algorithmic game theory / edited by Noam Nisan . . . [et al.]; foreword by Christos Papadimitriou.

p. cm.

Includes index.

ISBN-13: 978-0-521-87282-9 (hardback)

ISBN-10: 0-521-87282-0 (hardback)


1. Game theory. 2. Algorithms. I. Nisan, Noam. II. Title.

QA269.A43 2007

519.3-dc22 2007014231

ISBN 978-0-521-87282-9 hardback

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party Internet Web sites referred to in this publication and does not guarantee that any content on such Web sites is, or will remain, accurate or appropriate.



Algorithmic Game Theory

Over the last few years, there has been explosive growth in the research done at the interface of computer science, game theory, and economic theory, largely motivated by the emergence of the Internet. *Algorithmic Game Theory* develops the central ideas and results of this new and exciting area.

More than 40 of the top researchers in this field have written chapters whose topics range from the foundations to the state of the art. This book contains an extensive treatment of algorithms for equilibria in games and markets, computational auctions and mechanism design, and the “price of anarchy,” as well as applications in networks, peer-to-peer systems, security, information markets, and more.

This book will be of interest to students, researchers, and practitioners in theoretical computer science, economics, networking, artificial intelligence, operations research, and discrete mathematics.

Noam Nisan is a Professor in the Department of Computer Science at The Hebrew University of Jerusalem. His other books include *Communication Complexity*.

Tim Roughgarden is an Assistant Professor in the Department of Computer Science at Stanford University. His other books include *Selfish Routing and the Price of Anarchy*.

Éva Tardos is a Professor in the Department of Computer Science at Cornell University. Her other books include *Algorithm Design*.

Vijay V. Vazirani is a Professor in the College of Computing at the Georgia Institute of Technology. His other books include *Approximation Algorithms*.

Foreword

As the Second World War was coming to its end, John von Neumann, arguably the foremost mathematician of that time, was busy initiating two intellectual currents that would shape the rest of the twentieth century: game theory and algorithms. In 1944 (16 years after the minmax theorem) he published, with Oscar Morgenstern, his *Games and Economic Behavior*, thus founding not only game theory but also utility theory and microeconomics. Two years later he wrote his draft report on the EDVAC, inaugurating the era of the digital computer and its software *and its algorithms*. Von Neumann wrote in 1952 the first paper in which a polynomial algorithm was hailed as a meaningful advance. And, he was the recipient, shortly before his early death four years later, of Gödel's letter in which the P vs. NP question was first discussed.

Could von Neumann have anticipated that his twin creations would converge half a century later? He was certainly far ahead of his contemporaries in his conception of computation as something dynamic, ubiquitous, and enmeshed in society, almost organic – witness his self-reproducing automata, his fault-tolerant network design, and his prediction that computing technology will advance in lock-step with the economy (for which he had already postulated exponential growth in his 1937 Vienna Colloquium paper). But I doubt that von Neumann could have dreamed anything close to the Internet, the ubiquitous and quintessentially organic computational artifact that emerged after the end of the Cold War (a war, incidentally, of which von Neumann was an early soldier and possible casualty, and that was, fortunately, fought mostly with game theory and decided by technological superiority – essentially by algorithms – instead of the thermonuclear devices that were von Neumann's parting gift to humanity).

The Internet turned the tables on students of both markets and computation. It transformed, informed, and accelerated markets, while creating new and theretofore unimaginable kinds of markets – in addition to being itself, in important ways, a market. Algorithms became the natural environment and default platform of strategic decision making. On the other hand, the Internet was the first computational artifact that was not created by a single entity (engineer, design team, or company), but emerged from the strategic interaction of many. Computer scientists were for the first time faced with an object that they had to feel with the same bewildered awe with which economists have

always approached the market. And, quite predictably, they turned to game theory for inspiration – in the words of Scott Shenker, a pioneer of this way of thinking who has contributed to this volume, “the Internet is an equilibrium, we just have to identify the game.” A fascinating fusion of ideas from both fields – game theory and algorithms – came into being and was used productively in the effort to illuminate the mysteries of the Internet. It has come to be called algorithmic game theory.

The chapters of this book, a snapshot of algorithmic game theory at the approximate age of ten written by a galaxy of its leading researchers, succeed brilliantly, I think, in capturing the field’s excitement, breadth, accomplishment, and promise. The first few chapters recount the ways in which the new field has come to grips with perhaps the most fundamental cultural incongruity between algorithms and game theory: the latter predicts the agents’ equilibrium behavior typically with no regard to the ways in which such a state will be reached – a consideration that would be a computer scientist’s foremost concern. Hence, algorithms for computing equilibria (Nash and correlated equilibria in games, price equilibria for markets) have been one of algorithmic game theory’s earliest research goals. This body of work has become a valuable contribution to the debate in economics about the validity of behavior predictions: Efficient computability has emerged as a very desirable feature of such predictions, while computational intractability sheds a shadow of implausibility on a proposed equilibrium concept. Computational models that reflect the realities of the market and the Internet better than the von Neumann machine are of course at a premium – there are chapters in this book on learning algorithms as well as on distributed algorithmic mechanism design.

The algorithmic nature of mechanism design is even more immediate: This elegant and well-developed subarea of game theory deals with the design of games, with players who have unknown and private utilities, such that at the equilibrium of the designed game the designer’s goals are attained independently of the agents’ utilities (auctions are an important example here). This is obviously a computational problem, and in fact some of the classical results in this area had been subtly algorithmic, albeit with little regard to complexity considerations. Explicitly algorithmic work on mechanism design has, in recent years, transformed the field, especially in the case of auctions and cost sharing (for example, how to recover the cost of an Internet service from customers who value the service by amounts known only to them) and has become the arena of especially intense and productive cross-fertilization between game theory and algorithms; these problems and accomplishments are recounted in the book’s second part.

The third part of the book is dedicated to a line of investigation that has come to be called “the price of anarchy.” Selfish rational agents reach an equilibrium. The question arises: exactly how inefficient is this equilibrium in comparison to an idealized situation in which the agents would strive to collaborate selflessly with the common goal of minimizing total cost? The ratio of these quantities (the cost of an equilibrium over the optimum cost) has been estimated successfully in various Internet-related setups, and it is often found that “anarchy” is not nearly as expensive as one might have feared. For example, in one celebrated case related to routing with linear delays and explained in the “routing games” chapter, the overhead of anarchy is at most 33% over the optimum solution – in the context of the Internet such a ratio is rather insignificant

and quickly absorbed by its rapid growth. Viewed in the context of the historical development of research in algorithms, this line of investigation could be called “the third compromise.” The realization that optimization problems are intractable led us to approximation algorithms; the unavailability of information about the future, or the lack of coordination between distributed decision makers, brought us online algorithms; the price of anarchy is the result of one further obstacle: now the distributed decision makers have different objective functions. Incidentally, it is rather surprising that economists had not studied this aspect of strategic behavior before the advent of the Internet. One explanation may be that, for economists, the ideal optimum was never an available option; in contrast, computer scientists are still looking back with nostalgia to the good old days when artifacts and processes could be optimized exactly. Finally, the chapters on “additional topics” that conclude the book (e.g., on peer-to-peer systems and information markets) amply demonstrate the young area’s impressive breadth, reach, diversity, and scope.

Books – a glorious human tradition apparently spared by the advent of the Internet – have a way of marking and focusing a field, of accelerating its development. Seven years after the publication of *The Theory of Games*, Nash was proving his theorem on the existence of equilibria; only time will tell how this volume will sway the path of algorithmic game theory.

Paris, February 2007

Christos H. Papadimitriou

Preface

This book covers an area that straddles two fields, algorithms and game theory, and has applications in several others, including networking and artificial intelligence. Its text is pitched at a beginning graduate student in computer science – we hope that this makes the book accessible to readers across a wide range of areas.

We started this project with the belief that the time was ripe for a book that clearly develops some of the central ideas and results of algorithmic game theory – a book that can be used as a textbook for the variety of courses that were already being offered at many universities. We felt that the only way to produce a book of such breadth in a reasonable amount of time was to invite many experts from this area to contribute chapters to a comprehensive volume on the topic.

This book is partitioned into four parts: the first three parts are devoted to core areas, while the fourth covers a range of topics mostly focusing on applications. Chapter 1 serves as a preliminary chapter and it introduces basic game-theoretic definitions that are used throughout the book. The first chapters of Parts II and III provide introductions and preliminaries for the respective parts. The other chapters are largely independent of one another. The authors were requested to focus on a few results highlighting the main issues and techniques, rather than provide comprehensive surveys. Most of the chapters conclude with exercises suitable for classroom use and also identify promising directions for further research. We hope these features give the book the feel of a textbook and make it suitable for a wide range of courses.

You can view the entire book online at
www.cambridge.org/us/9780521872829
username: agt1user
password: camb2agt

Many people's efforts went into producing this book within a year and a half of its first conception. First and foremost, we thank the authors for their dedication and timeliness in writing their own chapters and for providing important

feedback on preliminary drafts of other chapters. Thanks to Christos Papadimitriou for his inspiring Foreword. We gratefully acknowledge the efforts of outside reviewers: Elliot Anshelevich, Nikhil Devanur, Matthew Jackson, Vahab Mirrokni, Herve Moulin, Neil Olver, Adrian Vetta, and several anonymous referees. Thanks to Cindy Robinson for her invaluable help with correcting the galley proofs. Finally, a big thanks to Lauren Cowles for her stellar advice throughout the production of this volume.

Noam Nisan
Tim Roughgarden
Éva Tardos
Vijay V. Vazirani

Contributors

Ross Anderson

Computer Laboratory
University of Cambridge

Moshe Babaioff

School of Information
University of California, Berkeley

Avrim Blum

Department of Computer Science
Carnegie Mellon University

Liad Blumrosen

Microsoft Research
Silicon Valley

John Chuang

School of Information
University of California, Berkeley

Bruno Codenotti

Istituto di Informatica e
Telematica, Consiglio
Nazionale delle Ricerche

Yevgeniy Dodis

Department of Computer Science
Courant Institute of Mathematical
Sciences, New York University

Joan Feigenbaum

Computer Science Department
Yale University

Michal Feldman

School of Business Administration
and the Center for the Study of Rationality
Hebrew University of Jerusalem

Eric Friedman

School of Operations Research
and Information Engineering
Cornell University

Jason D. Hartline

Microsoft Research
Silicon Valley

Kamal Jain

Microsoft Research
Redmond

Ramesh Johari

Department of Management Science
and Engineering
Stanford University

Anna R. Karlin

Department of Computer Science
and Engineering
University of Washington

Michael Kearns

Department of Computer
and Information Science
University of Pennsylvania

Jon Kleinberg

Department of Computer Science
Cornell University

Sébastien Lahaie

School of Engineering
and Applied Sciences
Harvard University

Ron Lavi

Faculty of Industrial Engineering
and Management, The Technion
Israel Institute of Technology

Mohammad Mahdian

Yahoo! Research
Silicon Valley

Yishay Mansour

School of Computer Science
Tel Aviv University

Tyler Moore

Computer Laboratory
University of Cambridge

Shishir Nagaraja

Computer Laboratory
University of Cambridge

Noam Nisan

School of Computer Science
and Engineering
Hebrew University of Jerusalem

Asuman Ozdaglar

Department of Electrical
Engineering and Computer
Science, MIT

Andy Ozment

Computer Laboratory
University of Cambridge

Christos H. Papadimitriou

Computer Science Division
University of California, Berkeley

David C. Parkes

School of Engineering
and Applied Sciences
Harvard University

David M. Pennock

Yahoo! Research
New York

Tal Rabin

T. J. Watson Research Center
IBM

Paul Resnick

School of Information
University of Michigan

Tim Roughgarden

Department of Computer Science
Stanford University

Amin Saberi

Department of Management
Science and Engineering
Stanford University

Rahul Sami

School of Information
University of Michigan

Michael Schapira

School of Computer Science
and Engineering
The Hebrew University of Jerusalem

James Schummer

M.E.D.S.
Kellogg School of Management
Northwestern University

Scott Shenker

EECS Department
University of California, Berkeley

R. Srikant

Department of Electrical and Computer
Engineering and Coordinated Science
Laboratory, University of Illinois at
Urbana-Champaign

Siddharth Suri

Department of Computer Science
Cornell University

Éva Tardos

Department of Computer Science
Cornell University

Kasturi Varadarajan

Department of Computer Science
University of Iowa

Vijay V. Vazirani

College of Computing
Georgia Institute of Technology

Berthold Vöcking

Department of Computer Science
RWTH Aachen University

Rakesh V. Vohra

M.E.D.S.
Kellogg School of Management
Northwestern University

Bernhard von Stengel

Department of Mathematics
London School of Economics

Tom Wexler

Department of Computer Science
Cornell University

Contents

<i>Foreword</i>	page xiii
<i>Preface</i>	xvii
<i>Contributors</i>	xix

I Computing in Games

1 Basic Solution Concepts and Computational Issues	3
<i>Éva Tardos and Vijay V. Vazirani</i>	
1.1 Games, Old and New	3
1.2 Games, Strategies, Costs, and Payoffs	9
1.3 Basic Solution Concepts	10
1.4 Finding Equilibria and Learning in Games	16
1.5 Refinement of Nash: Games with Turns and Subgame Perfect Equilibrium	18
1.6 Nash Equilibrium without Full Information: Bayesian Games	20
1.7 Cooperative Games	20
1.8 Markets and Their Algorithmic Issues	22
Acknowledgments	26
Bibliography	26
Exercises	26
2 The Complexity of Finding Nash Equilibria	29
<i>Christos H. Papadimitriou</i>	
2.1 Introduction	29
2.2 Is the NASH Equilibrium Problem NP-Complete?	31
2.3 The Lemke–Howson Algorithm	33
2.4 The Class PPAD	36
2.5 Succinct Representations of Games	39
2.6 The Reduction	41
2.7 Correlated Equilibria	45
2.8 Concluding Remarks	49
Acknowledgment	50
Bibliography	50

3 Equilibrium Computation for Two-Player Games in Strategic and Extensive Form	53
<i>Bernhard von Stengel</i>	
3.1 Introduction	53
3.2 Bimatrix Games and the Best Response Condition	54
3.3 Equilibria via Labeled Polytopes	57
3.4 The Lemke–Howson Algorithm	61
3.5 Integer Pivoting	63
3.6 Degenerate Games	65
3.7 Extensive Games and Their Strategic Form	66
3.8 Subgame Perfect Equilibria	68
3.9 Reduced Strategic Form	69
3.10 The Sequence Form	70
3.11 Computing Equilibria with the Sequence Form	73
3.12 Further Reading	75
3.13 Discussion and Open Problems	75
Bibliography	76
Exercises	77
4 Learning, Regret Minimization, and Equilibria	79
<i>Avrim Blum and Yishay Mansour</i>	
4.1 Introduction	79
4.2 Model and Preliminaries	81
4.3 External Regret Minimization	82
4.4 Regret Minimization and Game Theory	88
4.5 Generic Reduction from External to Swap Regret	92
4.6 The Partial Information Model	94
4.7 On Convergence of Regret-Minimizing Strategies to Nash Equilibrium in Routing Games	96
4.8 Notes	99
Bibliography	99
Exercises	101
5 Combinatorial Algorithms for Market Equilibria	103
<i>Vijay V. Vazirani</i>	
5.1 Introduction	103
5.2 Fisher’s Linear Case and the Eisenberg–Gale Convex Program	105
5.3 Checking If Given Prices Are Equilibrium Prices	108
5.4 Two Crucial Ingredients of the Algorithm	109
5.5 The Primal-Dual Schema in the Enhanced Setting	109
5.6 Tight Sets and the Invariant	111
5.7 Balanced Flows	111
5.8 The Main Algorithm	115
5.9 Finding Tight Sets	117
5.10 Running Time of the Algorithm	118
5.11 The Linear Case of the Arrow–Debreu Model	121
5.12 An Auction-Based Algorithm	122
5.13 Resource Allocation Markets	124

5.14	Algorithm for Single-Source Multiple-Sink Markets	126
5.15	Discussion and Open Problems	131
	Bibliography	132
	Exercises	133
6	Computation of Market Equilibria by Convex Programming	135
	<i>Bruno Codenotti and Kasturi Varadarajan</i>	
6.1	Introduction	135
6.2	Fisher Model with Homogeneous Consumers	141
6.3	Exchange Economies Satisfying WGS	142
6.4	Specific Utility Functions	148
6.5	Limitations	150
6.6	Models with Production	152
6.7	Bibliographic Notes	155
	Bibliography	156
	Exercises	158
7	Graphical Games	159
	<i>Michael Kearns</i>	
7.1	Introduction	159
7.2	Preliminaries	161
7.3	Computing Nash Equilibria in Tree Graphical Games	164
7.4	Graphical Games and Correlated Equilibria	169
7.5	Graphical Exchange Economies	176
7.6	Open Problems and Future Research	177
7.7	Bibliographic Notes	177
	Acknowledgments	179
	Bibliography	179
8	Cryptography and Game Theory	181
	<i>Yevgeniy Dodis and Tal Rabin</i>	
8.1	Cryptographic Notions and Settings	181
8.2	Game Theory Notions and Settings	187
8.3	Contrasting MPC and Games	189
8.4	Cryptographic Influences on Game Theory	191
8.5	Game Theoretic Influences on Cryptography	197
8.6	Conclusions	202
8.7	Notes	203
	Acknowledgments	204
	Bibliography	204
II Algorithmic Mechanism Design		
9	Introduction to Mechanism Design (for Computer Scientists)	209
	<i>Noam Nisan</i>	
9.1	Introduction	209
9.2	Social Choice	211
9.3	Mechanisms with Money	216
9.4	Implementation in Dominant Strategies	222

9.5	Characterizations of Incentive Compatible Mechanisms	225
9.6	Bayesian–Nash Implementation	233
9.7	Further Models	238
9.8	Notes	239
	Acknowledgments	240
	Bibliography	241
10	Mechanism Design without Money	243
	<i>James Schummer and Rakesh V. Vohra</i>	
10.1	Introduction	243
10.2	Single-Peaked Preferences over Policies	244
10.3	House Allocation Problem	253
10.4	Stable Matchings	255
10.5	Future Directions	262
10.6	Notes and References	263
	Bibliography	264
	Exercises	264
11	Combinatorial Auctions	267
	<i>Liad Blumrosen and Noam Nisan</i>	
11.1	Introduction	267
11.2	The Single-Minded Case	270
11.3	Walrasian Equilibrium and the LP Relaxation	275
11.4	Bidding Languages	279
11.5	Iterative Auctions: The Query Model	283
11.6	Communication Complexity	287
11.7	Ascending Auctions	289
11.8	Bibliographic Notes	295
	Acknowledgments	296
	Bibliography	296
	Exercises	298
12	Computationally Efficient Approximation Mechanisms	301
	<i>Ron Lavi</i>	
12.1	Introduction	301
12.2	Single-Dimensional Domains: Job Scheduling	303
12.3	Multidimensional Domains: Combinatorial Auctions	310
12.4	Impossibilities of Dominant Strategy Implementability	317
12.5	Alternative Solution Concepts	321
12.6	Bibliographic Notes	327
	Bibliography	327
	Exercises	328
13	Profit Maximization in Mechanism Design	331
	<i>Jason D. Hartline and Anna R. Karlin</i>	
13.1	Introduction	331
13.2	Bayesian Optimal Mechanism Design	335
13.3	Prior-Free Approximations to the Optimal Mechanism	339
13.4	Prior-Free Optimal Mechanism Design	344

13.5	Frugality	350
13.6	Conclusions and Other Research Directions	354
13.7	Notes	357
	Bibliography	358
	Exercises	360
14	Distributed Algorithmic Mechanism Design	363
	<i>Joan Feigenbaum, Michael Schapira, and Scott Shenker</i>	
14.1	Introduction	363
14.2	Two Examples of DAMD	366
14.3	Interdomain Routing	370
14.4	Conclusion and Open Problems	379
14.5	Notes	380
	Acknowledgments	381
	Bibliography	381
	Exercises	383
15	Cost Sharing	385
	<i>Kamal Jain and Mohammad Mahdian</i>	
15.1	Cooperative Games and Cost Sharing	385
15.2	Core of Cost-Sharing Games	387
15.3	Group-Strategyproof Mechanisms and Cross-Monotonic Cost-Sharing Schemes	391
15.4	Cost Sharing via the Primal-Dual Schema	394
15.5	Limitations of Cross-Monotonic Cost-Sharing Schemes	400
15.6	The Shapley Value and the Nash Bargaining Solution	402
15.7	Conclusion	405
15.8	Notes	406
	Acknowledgments	408
	Bibliography	408
	Exercises	410
16	Online Mechanisms	411
	<i>David C. Parkes</i>	
16.1	Introduction	411
16.2	Dynamic Environments and Online MD	413
16.3	Single-Valued Online Domains	417
16.4	Bayesian Implementation in Online Domains	431
16.5	Conclusions	435
16.6	Notes	436
	Acknowledgments	437
	Bibliography	437
	Exercises	439

III Quantifying the Inefficiency of Equilibria

17	Introduction to the Inefficiency of Equilibria	443
	<i>Tim Roughgarden and Éva Tardos</i>	
17.1	Introduction	443