The Knowledge Engineering Review, Vol. 26:1, 71–72.
doi:10.1017/S0269888910000287

Book Review

Algorithmic Game Theory by Noam Nisan, Tim Roughgarden, Éva Tardos and Vijay V. Vazirani, Cambridge University Press, 754 pp., £32.00, ISBN 0-521-87282-0

As we all know, over the past decade the world has become ever more connected as a result of the growth of the Internet and mobile communication. There is an ever-shrinking list of places where one is unable to connect, via telephone or data (or increasingly both), to the rest of the world. One less anticipated way in which this has changed the world is the impact it has had on computer science. With all this connectivity, computing devices now rarely operate in isolation, and the interaction between them is a subject of increasing importance. The importance of interaction is not just a function of the number of connected devices, but one that is also greatly affected by who owns them. If I own a million devices, I can program them to work together in a way that best suits me. If, however, I operate a network that connects a million devices operated by a million individuals, there are limits to my ability to enforce what I consider to be desirable behavior patterns on those million users of my network.

The desire to develop computer systems that can deal with many devices operated by individuals who are essentially autonomous has very naturally led computer scientists to look at game theory. Game theory, after all, is a well-developed set of mathematical techniques that aim to model the behavior of self-interested individuals, and *Algorithmic Game Theory* presents some excellent examples of work in the intersection of computer science and game theory, identifying a number of strands of work.

It starts with a group of papers under the heading 'Computing in Games'. These papers discuss one of the main contributions that has come out of the hybridization of computer science and game theory, that is, considering the computation of aspects of game theory. For example, one of the central ideas in game theory is that of equilibrium—that with enough knowledge about the possible actions of other individuals, and after suitable reflection, an individual can identify their best course of action in a way that factors in what works best for others. If all individuals do this, then we can predict how they will all behave and the system that comprises them all will operate at the equilibrium defined by the combination of all these 'best courses of action given what works best for others'. While this is an attractive concept, doing the 'suitable reflection' is far from trivial, and exactly how to compute different equilibria is the subject of many chapters in this section of the book.

The next section of the book is 'Algorithmic Mechanism Design'. Mechanism design is in some ways the flipside of the computation of equilibria. Whereas when computing equilibria we are interested in figuring out how individuals will act in some frameworks, mechanism design is about deciding how to construct a framework so that individuals behave in a desirable way. Of course, we still need to compute equilibria to determine the effectiveness of a specific framework, but the emphasis in mechanism design is on finding the best framework (mechanism) for a specific situation, and the chapters in this section of the book consider how to do this for a variety of applications.

The third section of the book 'Quantifying the Inefficiency of Equilibria' considers how the equilibrium behavior of systems—the behavior we get if we let all the individuals who use the system decide how to act—diverges from the behavior we would get if we could tell individuals how to behave. The issue here is that equilibria often involve individuals choosing an option that is less good for them than the one they would choose if they knew that other individuals were not going to try and exploit them—defending themselves from exploitation has a cost. The area in which this has been most studied to date is in routing traffic through networks, and this section of the book contains several chapters looking at different aspects of routing problems.

The final section of the book contains a number of miscellaneous topics, but ones that are no less interesting. For example, here we find an analysis of charging policies in peer-to-peer networks to alleviate the free-riding problem, and an analysis of prediction markets—essentially betting markets that, by careful choice of payoff, can be used to predict the likelihood of specific events and which have a good track record of correctly calling political elections.

Overall, this is a very good book. Like every collection, it suffers from the problem that it does not give uniform coverage to all the topics in a given area, and there is not as much introductory material or connective tissue between topics as one might wish for. However, that does not change the fact that the material it contains is of extremely high quality—there are some very good chapters—and that the price of the book makes it an absolute bargain.

Reviewed by Simon Parsons
Brooklyn College, City University of New York, USA