

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