CHAPTER 1

ALGORITHMIC GAME THEORY AND APPLICATIONS

Marios Mavronicolas
Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus & Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, 33102 Paderborn, Germany.

Vicky Papadopoulou
Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus.

Paul Spirakis
Department of Computer Engineering and Informatics, University of Patras, Rion, Patras 26500, Greece & Research Academic Computer Technology Institute, Rion, Patras 26500, Greece.

WILEY-
INTERSCIENCE

A JOHN WILEY & SONS, INC., PUBLICATION

 Algorithmic Game Theory and Applications. By M. Mavronicolas, V. Papadopoulou and P. Spirakis

CONTENTS

1.1 INTRODUCTION

Most of the existing and foreseen complex networks, such as the Internet, are operated and built by thousands of large and small entities (autonomous *agents*), which collaborate to process and deliver end-to-end *flows* originating from and terminating at any of them. The distributed nature of the Internet implies a lack of coordination among its users. Instead, each user attempts to obtain maximum performance according to his own parameters and objectives.

Methods from Game Theory and Mathematical Economics have been proven to be a powerful modeling tool, which can be applied to understand, control and efficiently design such dynamic, complex networks. Game Theory provides a good starting point for Computer Scientists in their endeavor to understand selfish rational behavior in complex networks with many *agents* (*players*). Such scenaria are readily modeled using techniques from Game Theory, where players with potentially conflicting goals participate in a common setting with well prescribed interactions.

Nash equilibrium [73, 74] distinguishes itself as the predominant concept of rationality in non-cooperative settings. So, Game Theory and its various concepts of equilibria provide a rich framework for modeling the behavior of selfish agents in these kinds of distributed or networked environments; they offer mechanisms to achieve efficient and desirable global outcomes in spite of the selfish behavior.

Mechanism Design, a subfield of Game Theory, asks how one can design systems so that agents' selfish behavior results to desired system-wide goals. *Algorithmic Mechanism Design* additionally considers computational tractability to the set of concerns of Mechanism Design. Work on Algorithmic Mechanism Design has focused on the complexity of centralized implementations of game-theoretic mechanisms for distributed optimization problems. Moreover, in such huge and heterogeneous networks, each agent does not have access to (and may not process) complete information. The notion of *bounded rationality* for agents, and the design of corresponding incomplete-information *distributed* algorithms, have been successfully utilized to capture the aspect of lack of global knowledge in information networks.

In this chapter, we review some of the most thrilling algorithmic problems and solutions, and corresponding advances, achieved on the account of Game Theory. The areas addressed are the followings:

Congestion Games. A central problem arising in the management of large-scale communication networks is that of routing traffic through the network. However, due to the large size of these networks, it is often impossible to employ a centralized traffic management. A natural assumption to make in the absence of central regulation is that network users behave selfishly and aim at optimizing their own individual welfare. One way to address this problem is to model this scenario as a non-cooperative multi-player game and formalize it using *congestion* game. Congestion games (either *unweighted* or *weighted*) offer a very natural framework for resource allocation in large networks like the Internet. In a nutshell, the main feature of congestion games is that they model *congestion* on a resource as a function of the number (or total weight) of all agents sharing the resource.

Price of Anarchy. We survey precise and approximate estimations for the *Price of Anarchy*; this is the cost of selfish behavior in dynamic, large-scale networks compared to hypothetical centralized solutions. We consider the Price of Anarchy for some of the most important network problems that are modeled by non-cooperative games; for example, we consider *routing* and *security* problems. A natural variant of the Price of Anarchy is the *Price of Stability* [5], which is the *best-case* cost of selfish behavior in complex networks, compared to a hypothetical centralized solution. The best-case assumption in the formulation of the Price of Stability implies that this cost can be enforced to the agents since they are interested in paying an as low cost as possible.

Selfish Routing with Incomplete Information. The impact of bounded rationality in networks with incomplete information can be addressed in two successful ways: by either *Bayesian games*, or by congestion games with *player specific payoff functions*. We will survey methods and tools for approximating network equilibria and network flows for a selfish system comprised of agents with bounded rationality.

Mechanism Design. *Mechanism Design* is a subfield of Game Theory and Microeconomics which deals with the design of protocols for rational agents. Generally, a Mechanism Design problem can be described as the task of selecting, out of a collection of feasible games, one which will yield desirable results for the designer. So, Mechanism Design can be thought of as the "inverse problem" in Game Theory, where the input is a game's outcome and the output is a game guaranteeing the desired outcome. The study of Mechanism Design from the algorithmic point of view starts with the seminal paper of Nisan and Ronen [76].

The routing problem in large-scale networks, where users are instinctively selfish, can be modeled by a non-cooperative game. Such a game could impose strategies that might induce an equilibrium close to the overall optimum. These strategies can be enforced through *pricing mechanisms* [28], *algorithmic mechanisms* [76] and *network design* [57, 87].

Stackelberg Games. We will examine network routing games from the *network designer*'s point of view. In particular, the network *administrator* or *designer* can define prices and rules, or even construct the network, in a way that induces near-optimal performance when the users act selfishly inside the system. Particularly interesting is the approach where the network manager takes part in the non-cooperative game. The manager has the ability to control centrally a part of the system resources, while the rest resources are managed by the selfish users. This approach has been implemented through *Stackelberg* or *Leader-Follower* games [16, 58].

The apparent advantage of this approach is that it might be easier to be deployed in large-scale networks. This is so since there is no need to add extra components to the network, or to exchange information between the users of the network.

In a typical Stackelberg game, one player acts as a *leader* (here, the centralized authority interested in optimizing system performance) and the rest act as *followers* (here, the selfish users). The problem is then to compute a strategy for the leader (a *Stackelberg strategy*) that induces the followers to react in a way that (at least approximately) minimizes the total latency in the system.

Selfish routing games can be modeled as a Stackelberg game. We will survey issues related to how the manager should assign the flow under his control into the system so as to induce optimal cost incurred by the selfish users. In particular, we will be interested in the complexity of designing *optimal* Stackelberg strategies.

Pricing mechanisms. *Pricing mechanisms* for resource allocation problems aim at allocating resources in such a way that those users who derive greater utility from the network are not denied access due to other users placing a lower value on it. In other words, pricing mechanisms are designed to guarantee *economic* efficiency. We will survey *cost-sharing* mechanisms for pricing the competitive usage of a collection of resources by a collection of selfish agents, each coming with an individual *demand*.

Network Security Games. We will also consider security problems in dynamic, large-scale, distributed networks. Such problems can be modeled as concise, non-cooperative multi-player games played on a graph. We will investigate the associated Nash equilibria for such network security games. In the literature, there have been studied at least two such interesting network security games.

Complexity of Computing Equilibria. The investigation of the computational complexity of finding a Nash equilibrium in a general strategic game is definitely a fundamental task for the development of Algorithmic Game Theory. Answers to such questions are expected to have great practical impact on both the analysis of the performance of antagonistic networks and on the development and implementation of policies for the network designers themselves.

Finding a Nash equilibrium in a game with two players could potentially be easier (than for many players) for several reasons.

- First, the *zero-sum* version of the game can be solved in polynomial time by linear programming. This grooms hopes for the polynomial solvability of the general (non-constant sum) version of the problem.
- Second, the two-players version of the game admits a polynomial size rational number solution, while there are games with three or more players that may only have solutions in irrational numbers.

This reasoning justified the identification of the problem of finding Nash equilibria for a 2-player game as one of the most important open questions in the field of Algorithmic Game Theory. The complexity of this problem was very recently settled in a perhaps surprising way in a series of breakthrough papers. In this chapter, we will later survey some of the worldwide literature related to this problem and the recent progress to it.

In this chapter we only assume a basic familiarity of the reader with some central concepts of Game Theory such as strategic games and Nash equilibria; for more details, we refer the interested reader to the leading textbooks [77, 78]. We also assume some acquaintance of the reader with the basic facts of the theory of computational complexity, as laid out, for example, in the leading textbook of Papadimitriou [80]. For readers interested in recalling the fundamental of algorithms design and analysis, we refer the reader to the prominent textbook of Kleinberg and Tardos [53]. For overwhelming motivation to dwelving into the secrets of Algorithmic Game Theory, we cheerfully refer the reader to the inspirational and prophetic survey of Papadimitriou in STOC 2001 [81].

1.2 CONGESTION GAMES

1.2.1 The General Framework

1.2.1.1 Congestion Games Rosenthal [84] introduced a special class of strategic games, now widely known as *congestion* games and currently under intense investigation by researchers in Algorithmic Game Theory. Here, the strategy set of each player is a subset of the power set of a set of *resources*; so, it is a set of sets of resources. Each player has an objective function, defined as the sum (over their chosen resources) of functions in the number of players sharing this resource. In his seminal work, Rosenthal showed with the help of a *potential function* that congestion games (in sharp contrast to *general* strategic games) always admit at least one pure Nash equilibrium.

An extension to congestion games are *weighted congestion games*, in which the players have *weights*, and thus exert different influences on the congestion of the resources. In (weighted) *network congestion games*, the strategy sets of the players correspond to paths in a network.

1.2.1.2 Price of Anarchy In order to measure the degradation of social welfare due to the selfish behavior of the players, Koutsoupias and Papadimitriou [60] introduced in their seminal work a global objective function, usually coined as *Social Cost*. It is quite remarkable that no notion similar in either spirit or structure to Social Cost had been studied in the Game Theory literature before. They defined the *Price of Anarchy*, also called *Coordination Ratio* and denoted as PoA, as the worst-case ratio between the value of Social Cost at a Nash equilibrium and that of some Social Optimum. The Social Optimum is the *best-case* Social Cost; so it is the least value of Social Cost achievable through cooperation. Thus, the Coordination Ratio measures the extent to which non-cooperation approximates cooperation.

As a starting point for analyzing the Price of Anarchy, Koutsoupias and Papadimitriou considered a very simple weighted network congestion game, now known as the *KP-model*. Here, the network consists of a single *source* and a single *destination* (in other words, it is a *single-commodity* network) which are connected together by parallel *links*. The *load* on a link is the total weight of players assigned to this link. Associated with each link is a *capacity* (or *speed*) representing the rate at which the link processes load. Each of the players selfishly routes from the source to the destination by using a probability distribution over the links. The private objective function of a player is its expected latency. The Social Cost is the expected maximum latency on a link, where the expectation is taken over all random choices of the players.

Fotakis *et al.* [34] have proved that computing Social Cost (in the form of expected maximum) is a $\#P$-complete problem. The stem of this negative result is the nature of exponential enumeration explicit in the definition of Social Cost (as an exponential-size expectation sum). An essentially identical $\#P$-hardness result has been proven recently by Daskalakis *et al.* [19]. This is one of the very few hard enumeration problems known in Algorithmic Game Theory as of today. Determining more remains a great challenge.

Mavronicolas and Spirakis [69] introduced *fully mixed Nash equilibria* for the particular case of the KP-model, in which each player chooses every link with positive probability. Gairing *et al.* [38, 39] explicitly conjectured that, in case the fully mixed Nash equilibrium exists, it is the worst-case Nash equilibrium with respect to Social Cost. This so-called *Fully Mixed Nash Equilibrium Conjecture* is simultaneously intuitive and significant.

- It is intuitive because the fully mixed Nash equilibrium favors an increased number of collisions between different players, since each player assigns its load with positive probability to every link. This increased probability of collisions should favor an increase to Social Cost.

- The conjecture is also significant since it identifies the worst-case Nash equilibrium over all instances. The Fully Mixed Nash Equilibrium Conjecture has been studied very intensively in the last few years over a variety of settings and models relative to the KP-model.

The KP-model was recently extended to *restricted strategy sets* [9, 35], where the strategy set of each player is a subset of the links. Furthermore, the KP-model was extended to general latency functions and studied with respect to different definitions of Social Cost [36, 37, 63].

Inspired by the arisen interest in the Price of Anarchy, the much older Wardrop model was reinvestigated in [88] (see also references therein). In this weighted network congestion game, weights can be split into arbitrary pieces. The social welfare of the system is defined as the sum of the edge latencies (*Sum* or *Total* Social Cost). An equilibrium in the Wardrop model can be interpreted as a Nash equilibrium in a game with infinitely many players, each carrying an infinitesimal amount of weight. There has been a tremendous amount of work following [88] on the reinvestigation of the Wardrop model. For an exposition, see the book by Roughgarden [86], which gives an account of the earliest results.

In [60], Koutsoupias and Papadimitriou initiated a systematic investigation of the social objective of (expected) maximum latency (also called *Maximum* Social Cost) for a weighted congestion game on uniformly related parallel links. The Price of Anarchy for this game has been shown to be $\Theta\left(\frac{\log m}{\log\log m}\right)$ if either the users or the links are identical [18, 59], and $\Theta\left(\frac{\log m}{\log\log\log m}\right)$ for weighted users and uniformly related links [18]. On the other hand, [17] shows that the Price of Anarchy is far worse and can be even unbounded for *arbitrary* latency functions. For uniformly related parallel links, identical users, and the objective of total latency, the Price of Anarchy is $1 - o(1)$ for the general case of mixed equilibria and $4/3$ for pure equilibria [63]. For identical users and *polynomial* latency functions of degree d, the Price of Anarchy is $d^{\Theta(d)}$ [8, 15].

Christodoulou and Koutsoupias [15] consider the Price of Anarchy of pure Nash equilibria in congestion games with linear latency functions. They showed that for general (asymmetric) games, the Price of Anarchy for Maximum Social Cost is $\Theta(\sqrt{n})$, where n is the number of players. For all other cases of symmetric or asymmetric games, and for both Maximum and Average Social Cost, the Price of Anarchy is shown to be $\frac{5}{2}$. Similar results were simultaneously obtained by Awerbuch *et al.* [15]

1.2.2 Pearls

A comprehensive survey of some of the most important recent advances in the literature on *atomic congestion games* is provided by [55]. That work is an overview of the extensive expertise on (mainly, network) congestion games and the closely related *potential games* [71], which has been developed in various disciplines (e.g., Economics, Computer Science and Operations Research) under a common formalization and modeling. In particular, the survey goes deep into the details of some of the most characteristic results in the area in order to compile a useful toolbox that Game Theory provides in order to study antagonistic behavior due to congestion phenomena in Computer Science settings.

1.2.2.1 Selfish Unsplittable Flows In [32], Fotakis *et al.* study congestion games where selfish users with varying service demands on the system resources may request a joint service from an arbitrary subset of resources. Each user's demand has to be served *unsplittably* from a specific subset of resources. In that work, it is proved that the weighted congestion games are no longer isomorphic to the well known potential games, although this was true for the case of users with identical service demands. The authors also demonstrate the power of the network structure in the case of users with varying demands. For very simple networks, they show that there may not exist a pure Nash equilibria, which is not true for the case of parallel links network or for the case of infinitely splittable service demands. Furthermore, the authors propose a family of networks (called *layered networks*) for which they show the existence of at least one pure Nash equilibrium when each resource charges its users with a delay equal to its load. Finally, the same work considers the Price of Anarchy for the family of layered networks in the same case. It is shown that the Price of Anarchy for this case is $\Theta\left(\frac{\log m}{\log\log m}\right)$. That is, within constant factors, the worst-case network is the simplest one (the *parallel links* network). This implies that, for this family of networks, the network structure does not affect the quality of the outcome of the congestion games played on the network in an essential way.

Panagopoulou and Spirakis [79] consider selfish routing in single-commodity networks, where selfish users select paths to route their loads (represented by arbitrary integer weights). They consider identical delay functions for the links of the network. That work focuses also on an algorithm suggested in [32]; this is a potential-based algorithm for finding pure Nash equilibria in such networks. The analysis of this algorithm from [32] has given an upper bound on its running time, which is polynomial in n (the number of users) and the sum W of their weights. This bound can be exponential in n when some weights are superpolynomial. Therefore, the algorithm is only known to be *pseudo–polynomial*. The work of Panagopoulou and Spirakis [79] provides strong experimental evidence that this algorithm actually converges to a pure Nash equilibria in polynomial time in n (and, therefore, independent of the weights values).

In addition, Panagopoulou and Spirakis [79] propose an initial allocation of users to paths that dramatically accelerates this algorithm, as opposed to an arbitrary initial allocation. A by–product of that work is the discovery of a weighted potential function when link loads are exponential to their loads. This guarantees the existence of pure Nash equilibria for these delay functions, while it extends the results of Fotakis *et al.* from [32].

1.2.2.2 Worst-Case Equilibria In [30], Fischer and Vöcking reexamined the question of worst-case Nash equilibria for the selfish routing game associated with the KP model [60], where n weighted jobs are allocated to m identical machines. Recall that Gairing *et al.* [38, 39] had conjectured that the fully mixed Nash equilibrium is the worst Nash equilibrium for this game (with respect to the expected maximum load over all machines). The known algorithms for approximating the Price of Anarchy relied on proven cases of that conjecture. In [30], the authors interestingly present a counter-example to the conjecture showing that fully mixed Nash equilibria cannot be generally used to approximate the Price of Anarchy within reasonable factors. In addition, they present an algorithm that constructs the so-called *concentrated Nash equilibria*, which approximate the worst-case Nash equilibrium within constant factors.

Although the work of Fischer and Vöcking [30] has disproved the Fully Mixed Nash Equilibrium Conjecture for the case of weighted users and identical links, the possibility that the conjecture holds for the case of identical users and arbitrary links is still open.

1.2.2.3 Symmetric Congestion Games Fotakis *et al.* [33] continued the work of [32] and studied computational and coordination issues of Nash equilibria in *symmetric* network congestion games. A game is *symmetric* if all users have the same strategy set and users costs are given by identical symmetric functions of other users' strategies. (Symmetric games were already considered in the original work of Nash [73, 74].) In unweighted congestions games, users are identical, so that a common strategy set implies symmetry.

This work proposed a simple and natural greedy method (which is called the *Greedy Best Response – GBR*), to compute a pure Nash equilibria. In this algorithm, each user plays only once and allocates his traffic to a path selected via a shortest path computation. It is shown that this algorithm works for three special cases: (1) *series-parallel networks*, (2) users are identical and (3) users are of varying demands but they have the same best response strategy for any initial network traffic (this is called the *Common Best Response* property).

The authors also give constructions where the algorithm fails if either the latter condition is violated (even for a series-parallel network), or the network is not series-parallel (even for the case of identical users). Thus, these results essentially indicate the limits of the applicability of this greedy approach.

The same work [33] studies also the Price of Anarchy for the objective of (expected) maximum latency. It is proved that for any network of m uniformly related links and for identical users, the Price of Anarchy is $\Theta\left(\frac{\log m}{\log\log m}\right)$. This result is complementary (and somewhat orthogonal) to a similar result proved in [32] for the case of weighted users to be routed in a layered network.

1.2.2.4 Exact Price of Anarchy Obtaining exact *bounds* on Price of Anarchy is, of course, the ultimate wish providing a happy end to the story. Unfortunately, the cases where such exact bounds are known are truly rare as of today. We describe here a particularly interesting example of a success story for one of these rare cases.

Exact bounds on the Price of Anarchy for both unweighted and weighted congestion games with polynomial latency functions are provided in [3]. The authors use the total latency as the Social Cost measure. The result in [3] vastly improve on results by Awerbuch *et al.* [8] and Christodoulou and Koutsoupias [15], where non-matching upper and lower bounds were given. (We will later discuss the precise relation of the newer result to the older results.)

For the case of *unweighted congestion games*, it is shown in [3] that the price of anarchy is exactly

$$\mathsf{PoA} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$ and Φ_d is a natural generalization of the *golden ratio* to larger dimensions such that Φ_d is the solution to the equation $(\Phi_d + 1)^d = \Phi_d^{d+1}$. The best known upper and lower bounds had before been shown to be of the form $d^{\,d(1-o(1))}$ [15]. However, the term $o(1)$ was still hiding a significant gap between the upper and the lower bound.

For *weighted congestion games*, the authors show that the Price of Anarchy is exactly

$$\mathsf{PoA} = \Phi_d^{d+1}.$$

This result closes the gap between the so far best upper and lower bounds of $O(2^{\,d} d^{d+1})$ and $\Omega(d^{d/2})$ from [8].

The authors of [3] show that the above values on the Price of Anarchy also hold for the subclasses of unweighted and weighted network congestion games. For the upper bounds, the authors use a similar analysis as in [15]. The core of their analysis is to simultaneously determine parameters c_1 and c_2 such that

$$y \cdot f(x+1) \leq c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y)$$

for all polynomial latency functions of maximum degree d and for all reals $x, y \geq 0$. For the case of unweighted users, it suffices to show the inequality for all pairs of integers x and y. (In order to prove their upper bound, Christodoulou and Koutsoupias [15] looked at the inequality with $c_1 = \frac{1}{2}$ and gave an asymptotic estimate for c_2.) In the analysis presented in [3], both parameters c_1 and c_2 are optimized. This optimization process required new mathematical ideas and is highly non-trivial. This optimization was successfully applied by Dumrauf and Gairing [24] to the so called *Polynomial Wardrop games*, where it yielded almost exact bounds on Price of Stability.

1.3 SELFISH ROUTING WITH INCOMPLETE INFORMATION

In his seminal work, Harsanyi [46] introduced an elegant approach to study non-cooperative games with *incomplete information*, where the players are uncertain about some parameters of the game. To model such games, he introduced the *Harsanyi transformation*, which converts a game with incomplete information to a strategic game where players may have different *types*. In the resulting *Bayesian game*, the players' uncertainty about each other's types is described by a probability distribution over all possible *type profiles*. It was only recently that Bayesian games were investigated from the point of view of Algorithmic Game Theory. Naturally, researchers were interested in formulating Bayesian versions of already studied routing games, as we described below.

In more detail, the problem of selfish routing with incomplete information has recently been faced via the introduction of new suitable models and the development of new methodologies that help to analyze such network settings. In particular, there were introduced new selfish routing games with incomplete information, called *Bayesian routing games* [40].

In a different piece of work, the same problem has been viewed as a congestion game where latency functions are *player-specific* [41], or a congestion game under the restriction that the link for each user must be chosen from a certain set of allowed links for the user [9, 26].

1.3.1 Bayesian Routing Games

Gairing *et al.* introduced [40] a particular selfish routing game with incomplete information, called *Bayesian routing game*. Here, n selfish *users* wish to assign their *traffics* to one of m parallel *links*. Users do not know each other's traffic. Following Harsanyi's approach, the authors introduce for each user a set of *types*. Each type represents a possible traffic; so, the set of types capture the set of all possibilities for each user. Unfortunately, users know the set of all possibilities for each other, but not the actual traffic itself.

Gairing *et al.* [40] proved, with the help of a potential function, that every Bayesian routing game has a pure Bayesian Nash equilibrium. This result has also been generalized to a larger class of games, called *weighted Bayesian congestion games*. For the case of identical links and *independent* type distributions, it is shown that a pure Bayesian Nash equilibrium can be computed in polynomial time. (A probability distribution over all possible type profiles is *independent* if it can be expressed as the product of independent probability distributions, one for each type.)

In the same work, Gairing *et al.* study structural properties of *Bayesian fully mixed Nash equilibria* for the case of identical links, they show that those maximize Individual Cost. This implies, in particular, that Bayesian fully mixed Nash equilibria maximize Social Cost as sum of Individual Costs.

In general, there may exist more than one fully mixed Bayesian Nash equilibrium. Gairing *et al.* [40] provide a characterization of the class of fully mixed Bayesian Nash equilibria for the case of independent type distribution; the characterization determines, in turn, the *dimension* of Bayesian fully mixed Nash equilibria. (The *dimension* of Bayesian fully mixed Nash equilibria is the dimension of the smallest Euclidean space into which all Bayesian fully mixed Nash equilibria can be mapped.)

Finally, Gairing *et al.* consider [40] the Price of Anarchy for the case of identical links and for three different Social Cost measures; that is, they consider Social Cost as Expected Maximum Congestion, as Sum of Individual Costs and as Maximum Individual Cost. For the latter two measures, (asymptotic) tight bounds were provided using the proven structural properties of fully mixed Bayesian Nash equilibria.

1.3.2 Player-Specific Latency Functions

Gairing *et al.* [41] address the impact of incomplete knowledge in (weighted) network congestion games with either splittable or unsplittable flow. In this perspective, the proposed models generalize the two famous models of selfish routing, namely weighted (network) congestion games and Wardrop games, to accommodate player-specific latency functions. Latency functions may be arbitrary, non-decreasing functions; however, many of the shown results in [41] assume that the latency function for player i on resource j is a *linear* function $f_{ij}(x) = a_{ij}x + b_{ij}$, where $a_{ij} \geq 0$ and $b_{ij} \geq 0$. Gairing *et al.* use the term *player-specific capacities* to denote a game where $b_{ij} = 0$ in all (linear) latency functions.

Gairing *et al.* [41] derive several interesting results on the existence and computational complexity of (pure) Nash equilibria and on the Price of Anarchy. For routing games on parallel links with player-specific capacities, they introduce two new *potential* functions, one for unsplittable and for splittable traffics. The first potential function is used to prove that games with unweighted players possess the *finite improvement property* in the case of unsplittable traffics. It is also shown in [41] that games with weighted players do not possess the finite improvement property in general, even if there are only three users. The second potential function is a convex function tailored to the case of splittable traffics. This convex function is minimized if and only if the corresponding assignment is a Nash equilibrium. Since such minimization of a convex latency function can be carried out in polynomial time, the established equivalence between minimizes of the potential function and Nash equilibria implies that a Nash equilibrium can be computed in polynomial time.

The same work [41] proves upper and lower bounds on the Price of Anarchy under a certain restriction on the linear latency functions. For the case of unsplittable traffics, the upper and lower bounds are asymptotically tight. All bounds on the Price of Anarchy translate to corresponding bounds for general congestion games.

1.3.3 Network Uncertainty in Selfish Routing

The problem of selfish routing in the presence of incomplete network information has also been studied by Georgiou *et al.* [43]. This work proposes an interesting new model for selfish routing in the presence of incomplete network information. The model proposed by Georgiou *et al.* captures situations where the users have incomplete information regarding the link capacities. Such uncertainty may be caused if the network links actually represent complex paths created by *routers*, which are constructed differently on separate occasions and sometimes according to the presence of congestion or link failures.

The new, extremely interesting model presented in [43] consists of a number of users who wish to route their traffic on a network of m parallel links with the objective of minimizing their latency. In order to capture the lack of precise knowledge about the capacity of the network links, Georgiou *et al.* [43] assumed that links may present a number of different capacities. Each user's uncertainty about the capacity of each link is modeled via a probability distribution over all possibilities. Furthermore, it is assumed that users may have different sources of information regarding the network; therefore, Georgiou *et al.* assume the probability distributions of the various users to be (possibly) distinct from each other. This gives rise to a very interesting model with user-specific payoff functions, where each user uses its distinct probability distribution to take decisions as to how to route its traffic.

The authors propose simple polynomial time algorithms to compute pure Nash equilibria in some special cases of the problem and demonstrate that a counter-example presented in [70], showing that pure Nash equilibria may not exist in the general case, does not apply to their model. Thus, Georgiou *et al.* identify an interesting open problem in this area, that of the existence of pure Nash equilibria in the general case of their model. Also, two different expressions for the Social Cost and the associated Price of Anarchy are identified and employed in [43]. For the latter, Georgiou *et al.* obtain upper bounds for the general case and some better upper bounds for several special cases of their model.

In the same work, Georgiou *et al.* show how to compute the fully mixed Nash equilibrium in polynomial time; they also show that when it exists, it is unique. Also, Georgiou *et al.* prove that for certain instances of the game, fully mixed Nash equilibria assign all links to all users equiprobably. Finally, the work in [43] verifies the Fully Mixed Nash Equilibrium conjecture, namely that the fully mixed Nash equilibrium maximizes Social Cost.

1.3.4 Restricted Selfish Scheduling

Elsässer *et al.* [26] further consider selfish routing problems in networks under the restriction that the link for each user must be chosen from a certain set of allowed links for the user. It is particularly assumed that each user has access (that is, finite cost) to only *two* machines; its cost on other machines is infinitely large, giving it no incentive to switch there. Interaction with just a few neighbors is a basic design principle to guarantee efficient use of resources in a distributed system. Restricting the number of interacting neighbors to just two is then a natural starting point for the theoretical study of the impact of selfish behavior in a distributed system with local interactions. In the model of Elsässer *et al.*, the (expected) cost of a user is the (expected) load on the machine it chooses.

The particular way of modeling local interaction in [26] has given rise to a simple, graph-theoretic model for selfish *scheduling* among m non-cooperative *users* over a collection of n *machines* with local interaction. In their graph-theoretic model, Elsässer *et al.* [26] address these bounded interactions by using an *interaction graph*, whose vertices and edges are the machines and the users, respectively. Elsässer *et al.* [26] have been interested in the impact of their modeling on the properties of the induced Nash equilibria.

The main result of Elsässer *et al.* [26] is that the *parallel links* graph is the *best-case* interaction graph – the one that minimizes expected *makespan* of the *standard fully mixed Nash equilibrium* – among all *3-regular* interaction graphs. (In the standard fully mixed Nash equilibria each user chooses each of its two admissible machines with probability $\frac{1}{2}$). The proof employs a graph-theoretic lemma about *orientations* in 3-regular graphs, which may be of independent interest. This is a particularly pleasing case where Algorithmic

Game Theory rewards Graph Theory with a wealth of new interesting problems about orientations in regular graphs.

A lower bound on Price of Anarchy is also provided in the work of Elsässer *et al.* [26]. In particular, it is proved that there is an interaction graph incurring Price of Anarchy $\Omega\left(\frac{\log n}{\log \log n}\right)$. This bound relies on a proof employing pure Nash equilibria. Finally, the authors present counterexample interaction graphs to prove that a *fully mixed Nash equilibrium* may sometimes not exist at all. (A characterization of interaction graphs admitting fully mixed Nash equilibria is still missing.) Moreover, they prove existence and uniqueness properties of the fully mixed Nash equilibrium for *complete bipartite* graphs and *hypercube* graphs.

The problems left open in [26] invite Graph Theory to a pleasing excursion into Algorithmic Game Theory.

1.3.5 Adaptive Routing with Stale Information

Fischer and Vöcking [29] consider the problem of adaptive routing in networks by selfish users that lack central control. The main focus of this work is on simple adaption policies, or *dynamics*, that make possible use of stale information. The analysis provided in [29] covers a wide class of dynamics encompassing the well-known *replicator dynamics* and other dynamics from *Evolutionary Game Theory*; the basic milestone is the well known fact that choosing the best option on the basis of out-of-date information can lead to undesirable oscillation effects and poor overall performance.

Fischer and Vöcking [29] show that it is possible to cope with this problem, and guarantee efficient convergence towards an equilibrium state, for all of this broad class of dynamics, if the function describing the cost of an edge depending on its load is not too steep. As it turns out, guaranteeing convergence depends solely on the size of a single parameter describing the greediness of the agents!

While the best response dynamics, which corresponds to always choosing the best option, performs well if information is always up-to-date, it is interestingly clear from the results in [29] that this policy fails when information is stale. More interestingly, Fischer and Vöcking [29] present a dynamics which approaches the global optimal solution in networks of parallel links with linear latency functions as fast as the best response dynamics does, but which does not suffer from poor performance when information is out-of-date.

1.4 ALGORITHMIC MECHANISM DESIGN

Mechanism Design is a subfield of Game Theory and Microeconomics which, generally speaking, deals with the design of protocols for rational agents. In most simple words, a Mechanism Design problem can be described as the task of selecting from a collection of (feasible) games, a game which will yield desirable results for the designer. Specifically, the theory of Mechanism Design has focused on problems where the goal is to satisfactorily aggregate privately known preferences of several agents towards a *social choice*. Intuitively, a Mechanism Design problem has two components:

- The usual algorithmic output specification
- Descriptions of what the participating agents want, formally given as *utility functions* over the set of possible *outputs* (outcomes).

The origin of Algorithmic Mechanism Design is marked with the seminal paper of Nisan and Romen [76].

A mechanism solves a given problem by assuring that the required outcome occurs, under the assumption that agents choose their strategies as to maximize their own selfish utilities. A mechanism needs thus to ensure that players' utilities (which it can influence by handing out *payments*) are compatible with the algorithm.

Recall that the routing problem in large-scale networks where users are instinctively selfish can be modeled as a non-cooperative game. Such a game is expected to impose strategies that would induce an equilibrium as close to the overall optimum as possible. Two possible approach to formulate such strategies are through *pricing mechanisms* [28] and *network design* [57, 87].

In the first approach, the network administrator defines *prices* (or *rules*) in a way that induces near optimal performance when the users act selfishly. This approach has been considered in [10, 16] (see also references therein). In the second approach, the network manager takes part in the noncooperative game. The manager has the ability to control centrally a part of the system resources, while the rest of the resources are to be shared by the selfish users. This approach has been studied through *Stackelberg* or *Leader-Follower* games [50, 85] (see also references therein). We here overview some issues related to how should the manager assign the flow he controls into the system, with the objective to induce optimal cost in spite of the behavior of the selfish users.

1.4.1 Stackelberg Games

In [85], Roughgarden studies the problem of optimizing the performance of a system shared by selfish, noncooperative users assigned to shared machines with load-dependent latency functions. Roughgarden measures system performance by the total latency of the system. (This measure is different than that used in the KP-model.) Assigning jobs according to the selfish interests of individual users typically results in suboptimal system performance. However, in many systems of this type, there is a mixture of "selfishly controlled" and "centrally controlled" jobs; as the assignment of centrally controlled jobs will influence the subsequent actions by selfish users, the degradation in system performance due to selfish behavior can be reduced by scheduling the centrally controlled jobs in the best possible way. Stackelberg games provide a framework that fits this situation in an excellent way.

A *Stackelberg game* is a special game where there are two kinds of entities: a number of selfish entities, called *players*, that are interested in optimizing their own utilities, and a distinguished *leader* controlling a number of non-self-interested entities called *followers*; the leader aims at improving the social welfare and decides on the strategies of the followers so that the resulting situation will induce suitable decisions for the players that will optimize social welfare (as much as possible).

Roughgarden [85] formulates this particular goal for such a selfish routing system as an optimization problem via Stackelberg games. The problem is then to compute a strategy for the leader (a *Stackelberg strategy*) that induces the followers to react in a way that (at least approximately) minimizes the total latency in the system. Roughgarden [85] proves that, perhaps not surprisingly, it is $\mathcal{NP}$-hard to compute the *optimal* Stackelberg strategy; he also presents simple strategies with provable performance guarantees.

More precisely, Roughgarden [85] gives a simple algorithm to compute a strategy inducing a job assignment with total latency no more than a small constant times that of the optimal assignment for all jobs; in the absence of centrally controlled jobs and a Stackelberg strategy, no result of this type is possible. Roughgarden also proves stronger performance guarantees in the special case where every latency function is linear in the load.

1.4.1.1 The Price of Optimum Kaporis and Spirakis continued in [50] the study of the Stackelberg games from [85]. They considered a system of parallel machines, each with a strictly increasing and differentiable load dependent latency function. The users of such a system are of infinite number and act selfishly, routing their infinitesimally small portion of the total flow they control to machines of currently minimum delay. In that work, such a system is modeled as a Stackelberg or Leader-Followers game motivated by [88].

In [85], Roughgarden had presented the `LLF` Stackelberg strategy for a *Leader* in a Stackelberg game with an infinite number of *Followers*, each routing its infinitesimal flow through machines of currently minimum delay (this is called the *Flow Model* in [85]). An important question posed there was the computation of the *least* portion β_M that a Leader must control in order to enforce the overall Optimum Cost on the system. In [50], an algorithm that computes β_M was presented and its optimality was also shown. Most importantly, it was proved that the algorithm presented is *optimal* for *any* class of latency functions for which Nash and optimum assignments can be efficiently computed. This is one of a very few known cases where the computation of optimal Stackelberg strategies is reduced to the computation of (pure) Nash equilibria and optimal assignments.

1.4.2 Cost Sharing Mechanisms

In its most general form, a *Cost Sharing Mechanism* specifies how costs originating from resource consumption in a selfish system should be shared among the users of the system. Apparently, not all sharing ways are good. Intuitively, a cost sharing mechanism is good if it can induce equilibria optimizing social welfare as much as possible. This point of view was adopted in a recent work by Mavronicolas *et al.* [65].

In more detail, a simple and intuitive *cost mechanism* which assigns *costs* for the competitive usage of m *resources* by n selfish *agents* was proposed by Mavronicolas *et al.* [65]. Each agent has an individual *demand*; demands are drawn according to some (unknown) probability distribution coming from a (known) class of probability distributions. The cost paid by an agent for a resource he chooses is the total demand put on the resource divided by the number of agents who chose that same resource. So, resources charge costs in an equitable, fair way, while each resource makes no *profit* out of the agents. [1]This simple model was called *Fair Pricing* in [65].

Mavronicolas *et al.* analyzed in [65] the *Nash equilibria* (both *pure* and *mixed*) for the induced game; in particular, they consider the *fully mixed Nash equilibrium*, where each agent selects each resource with non-zero probability. While offering (in addition) an advantage with respect to convenience in handling, the fully mixed Nash equilibrium is suitable for that economic framework under the very natural assumption that each resource offers usage to all agents without imposing any access restrictions.

The most significant contribution of [65] was the introduction of the *Diffuse Price of Anarchy* for the analysis of Nash equilibria in the induced game. Roughly speaking, the Diffuse Price of Anarchy is an extension to the Price of Anarchy that takes into account the probability distribution of the demands. Roughly speaking, the Diffuse Price of Anarchy is

[1]One could argue that this pricing scheme is *unfair* in the sense that players with smaller demands can be forced to support those players with larger demands that share the same resource. However, the model can also coined as fair on account of the fact that it treats all players sharing the same resource equally, and players are not overcharged beyond the actual cost of the resource they choose

the *worst-case*, over all allowed probability distributions, of the expectation (according to each specific probability distribution) of the ratio of Social Cost over Optimum in the *worst-case* Nash equilibrium. The Diffuse Price of Anarchy is meant to alleviate the sometimes overly pessimistic Price of Anarchy due to Koutsoupias and Papadimitriou [60] (which is a *worst-case* measure) by introducing and analyzing stochastic assumptions on the system inputs.

Mavronicolas *et al.* [65] proved that pure Nash equilibria may not exist unless all chosen demands are identical; in contrast, a fully mixed Nash equilibrium exists for all possible choices of the demands. Further on, it was proved that the fully mixed Nash equilibrium is the *unique* Nash equilibrium in case there are only two agents. It was also shown that, in the *worst-case* choice of demands, the Price of Anarchy is $\Theta(n)$; for the special case of two agents, the Price of Anarchy is less than $2 - \frac{1}{m}$.

A plausible assumption is that demands are drawn from a *bounded, independent probability distribution*, where all demands are *identically distributed* and each is at most a (*universal* for the class) constant times its expectation. Under this very general assumption, it is proved in [65] that the Diffuse Price of Anarchy is at most that same universal constant; the constant is just 2 when each demand is distributed symmetrically around its expectation.

1.4.3 Tax Mechanisms

How much can *taxes* improve the performance of a selfish system? This is a very general question since it leaves three important dimensions of it completely unspecified: the precise way of modeling taxes, the selfish system itself, and the measure of performance. Making specific choices for these three dimensions gives rise to specific interesting questions about taxes. There is already a sizeable amount of literature addressing such questions and variants of them (see, for example, [10, 16, 31] and references therein). In this section, we briefly describe the work of Caragiannis *et al.* [10], and we refer the reader to [16, 31] for additional related results.

Caragiannis *et al.* [10] consider the (by now familiar) class of congestion games due to Rosenthal [84] as their selfish system; they consider several measures for social welfare, including total latency and a new interesting measure they introduce, called *total disuitility*, which is the sum of latencies plus taxes incurred to players. Caragiannis *et al.* [10] focus on the well studied case of linear latency functions, and they provide many (both positive and negative) interesting results.

Their most interesting positive result is (in our opinion) the fact that there is a way to assign taxes that can improve the performance of congestion games by forcing players to follow strategies by which the total latency is within a factor of 2 of the least possible; Caragiannis *et al.* prove that, most interestingly, this is the *best* possible way of assigning taxes. Furthermore, Caraginannis *et al.* [10] consider cases where the system performance may be very poor in the absence of taxes; they prove that, fortunately, in such cases the total disuitility *cannot* be much larger than the *optimal* total latency. Another interesting result eminating from the work of Caragiannis *et al.* [10] is that there is a polynomial time algorithm (based on solving convex quadratic programs) to compute good taxes; this represents the *first* result on the efficiency of taxes for linear congestion games.

1.5 NETWORK SECURITY GAMES

It is an undeniable fact that the huge growth of the Internet has significantly extended the importance of *Network Security* [90]. Unfortunately, as it is well known, many widely used Internet systems and components are prone to security risks (see, for example, [14]); some of these risks have even led to successful and well-publicized attacks [89]. Typically, an *attack* exploits the discovery of loopholes in the security mechanisms of the Internet. Attacks and *defenses* are currently attracting a lot of interest in major forums of communication research. A current challenge for Algorithmic Game Theory is to invent and analyze appropriate theoretical models of security attacks and defenses for emerging networks like the Internet.

Two independent research teams, one consisting of Aspnes *et al.* [6] and another consisting of Mavronicolas *et al.* [67, 68], initiated recently the introduction of strategic games on graphs (and the study of their associated Nash equilibria) as a means of studying security problems in networks with selfish entities. The non-trivial results achieved by these two teams exhibit a novel interaction of ideas, arguments and techniques from two seemingly diverse fields, namely *Game Theory* and *Graph Theory*. This research line invites a simultaneously game-theoretic and graph-theoretic analysis of network security problems, where not only threats seek to maximize their caused damage to the network, but also the network seeks to protect itself as much as possible.

The two graph-theoretic models of Internet security can be cast as particular cases of the so called *Interdependent Security* games studied earlier by Kearns and Ortiz [52]. There, a large number of players must make individual decisions related to security. The ultimate safety of each player may depend in a complex way on the actions of the entire population.

1.5.1 A Virus Inoculation Game

Aspnes *et al.* [6] consider an interesting graph-theoretic game with an interesting security flavor, modeling containment of the spread of *viruses* on a network with installable *antivirus* software. In this game, the antivirus software may be installed at individual nodes; a virus damages a node if it can reach the node starting at a random initial node and proceeding to it without crossing a node with installed antivirus software. Aspnes *et al.* [6] prove several algorithmic properties for their graph-theoretic game and establish connections to a certain graph-theoretic problem called *Sum-of-Squares Partition*.

Moscibroda *et al.* [72] initiate the study of *Byzantine Game Theory* in the context of the specific virus inoculation game introduced by Aspnes *et al.* [6]. In their extention, they allow some players to be malicious or *Byzantine* rather than selfish. They ask the very natural question of what the impact of Byzantine players on the performance of the system compared to either the purely selfish setting (where all players are self-interested and there are no Byzantine players) or to the social optimum is.

To address such questions, they introduce the very interesting notion of the *Price of Malice* which captures the efficiency degradation due to the presence of Byzantine players (on top of selfish players). Moscibroda *et al.* [72] use the Price of Malice to quantify how much the presence of Byzantine players can deteriorate the social welfare of the distributed system corresponding to the virus inoculation game of Aspnes *et al.* [6]. Most interestingly, Moscibroda *et al.* [72] demonstrate that in case the selfish players are highly *risk-averse*, the social welfare of the system can improve as a result of taking Byzantine players into account!

We expect that Byzantine Game Theory will further develop in the upcoming years and be applied successfully to evaluate the impact of Byzantine players on the performance of selfish computer systems.

1.5.2 A Network Security Game

The work of Mavronicolas *et al.* [67, 68] considers a security problem on a distributed network modeled as a multi-player non-cooperative game with *attackers* (e.g., viruses) and a *defender* (e.g., a security software) entities. More specifically, there are two classes of confronting randomized players on a graph: ν *attackers*, each choosing vertices and wishing to minimize the probability of being caught, and a single *defender*, who chooses edges and gains the expected number of attackers it catches. The authors exploit both game-theoretic and graph-theoretic tools for analyzing the associated Nash equilibria.

In a subsequent work, Mavronicolas *et al.* [64] introduced the *Price of Defense* in order to evaluate the loss in the provided security guarantees due to the selfish nature of attacks and defenses. The work address the question of whether there are Nash equilibria that both are computationally tractable and offer good Price of Defense. An extensive collection of trade-offs between Price of Defense and the computational complexity of Nash equilibria is provided in the work of Mavronicolas *et al.* [64]. Most interestingly, the work of Mavronicolas *et al.* [64, 66, 67, 68] introduce certain natural classes of Nash equilibria for their network security game on graphs, including *Matching Nash equilibria* [67, 68] and *Perfect Matching Nash equilibria* [64]; they prove that deciding the existence of equilibria from such classes is precisely equivalent to the recognition problem for *König-Egervary* graphs [25, 54]. So, this establishes a very interesting (and perhaps unexpected) link to some classical pearls in Graph Theory.

1.6 COMPLEXITY OF COMPUTING EQUILIBRIA

By Nash's celebrating result [73, 74] every strategic game has at least one Nash equilibrium (and an odd number of them). What is the complexity of computing one? Note that this question is meaningful exactly when the payoff table is given in some implicit way that allows for a succinct representation. The celebrated algorithm of Lemke and Howson [61] shows that for bimatrix games this complexity is no more than exponential.

1.6.1 Pure Nash Equilibria

A core question in the study of Nash equilibria is which games have pure Nash equilibria. Also, under what circumstances can we find one (assuming that there is one) in polynomial time?

Recall that congestion games make a class of games that are guaranteed to have pure Nash equilibria. In a classical paper [84], Rosenthal proves that, in any such game, the *Nash dynamics* converges; equivalently, the directed graph with action combinations as nodes and payoff–improving deviations by individual players as edges is acyclic. Hence, the game has pure Nash equilibria which are the *sinks* of this graph. The proof is based on a simple potential function. This existence theorem, however, again left open the question of whether there is a polynomial-time algorithm for finding pure Nash equilibria in congestion games.

Fabrikant *et al.* [27] prove that the answer to this general question is positive when all players have the same origin and destination (the so-called *symmetric* case); a pure Nash equilibrium is found by computing the optimum of Rosenthal's potential function through a reduction to *min-cost flow*. However, it is shown that computing a pure Nash equilibrium in the general network case is $\mathcal{PLS}$-complete [49]. Intuitively, this means that it is as hard to compute as any object whose existence is guaranteed by a potential function. (The precise definition of the complexity class $\mathcal{PLS}$ is beyond the scope of this chapter.) The proof of [27] has the interesting consequence the existence of examples with exponentially long shortest paths, as well as the $\mathcal{PSPACE}$-completeness for the problem of computing a Nash equilibrium reachable from a specified state.

The completeness proof requires reworking the reduction to the problem of finding local optimal of weighted MAX2SAT instances. Ackermann *et al.* [1] present a significantly simpler proof based on a $\mathcal{PLS}$-reduction from MAX-CUT showing that finding Nash equilibria in network congestion games is $\mathcal{PLS}$-complete even for the case of linear latency functions. Additional results about the complexity of pure Nash equilibria in congestion games appear in the works of Ackermann *et al.* [1, 2].

Gottlob *et al.* [45] provide a comprehensive study of complexity issues related to pure Nash equilibria. They consider restrictions of strategic games intended to capture certain aspects of bounded rationality. For example, they show that even in the settings where each player's payoff function depends on the strategies of at most three other players, and where each player is allowed to choose one out of at most three strategies, the problem of determining whether a game has a pure Nash equilibrium is $\mathcal{NP}$-complete. On the positive side, they also identified tractable classes of games.

1.6.2 Mixed Nash Equilibria

Daskalakis *et al.* [20] consider the complexity of Nash equilibria in a game with four or more players. They show that this problem is complete for the complexity class $\mathcal{PPAD}$. Intuitively, this means that a polynomial-time algorithm would imply a similar algorithm, e.g., for computing *Brouwer fixpoints*; note that this is a problem for which quite strong lower bounds for large classes of algorithms are known [48]. (A precise definition of the complexity class $\mathcal{PPAD}$ is beyond the scope of this chapter.)

Nash [73, 74] had shown his celebrated result on the existence of Nash equilibria by reducing the existence of Nash equilibria to the existence of Brouwer fixpoints. Given any strategic game, Nash constructs a Brouwer function whose fixpoints are precisely the equilibria of the game. In Nash's reduction, as well as in subsequent simplified ones [42], the constructed Brouwer function is quite specialized; this has led to the speculation that the fixpoints of such functions (thus, Nash equilibria) are easier to find than for *general* Brouwer functions. In [20], this question is answered in the negative by presenting a very interesting reduction in the opposite direction: Any (computationally presented) Brouwer function can be simulated by a suitable game, so that Nash equilibria correspond to fixpoints.

In [23], it is proved that computing a Nash equilibrium in a 3-player game is also $\mathcal{PPAD}$-complete. The proof is based on a variant of an *arithmetical gadget* from [44], Independently, Chen and Deng [11] have also come up with a quite different proof of the same result.

In a very recent paper [12], Chen and Deng settle the complexity of Nash equilibria for 2-player strategic games with a $\mathcal{PPAD}$-completeness proof. Their proof derived a direct reduction from a search problem called the *3-Dimensional Brouwer* problem, which is known to be $\mathcal{PPAD}$-complete [20] to the objective problem. The completeness proof of [12] utilizes new gadgets for various arithmetic and logic operations.

1.6.3 Approximate Nash Equilibria

As it is always the case, an established intractability invites an understanding of the limits of approximation. Since it was established thta computing a Nash equilibrium is $\mathcal{PPAD}$-complete [20], even for 2-players strategic games [12], the question of computing approximate Nash equilibria has emerged as the central remaining open problem in the area of computing Nash equilibria.

Assume from this point on that all utilities have been normalized to be between 0 and 1. (Clearly, this assumption is without any loss of generality.) Say that a set of mixed strategies is an *ε-approximate Nash equilibrium*, where $\varepsilon > 0$, if for each player all strategies have expected payoff that is at most ε more that the expected payoff for its strategy in the given set. (So, ε is an additive approximation term.)

Lipton *et al.* [62] proved that an ε-approximate Nash equilibrium can be computed in time $O(n^{\frac{\log n}{\epsilon^2}})$ (that is, in strictly subexponential time) by examining all supports of size $\frac{\log n}{\epsilon^2}$. It had been earlier pointed out [4] that no algorithm examining supports smaller than about $\log n$ can achieve an approximation better than $\frac{1}{4}$, even for zero-sum games. In addition, it is easy to see that a $\frac{3}{4}$-approximation Nash equilibrium can be found (in polynomial time) by examining all supports of size 2.

Two research teams, one consisting of Daskalakis *et al.* [21] and the other of Kontogiannis *et al.* [56] investigated very recently the approximability of Nash equilibria in 2-player games, and established essentially identical, strong results. Most remarkably, there is a simple, linear-time algorithm in [21], which builds heavily on a corresponding algorithm from [56]; it examines just two strategies per player and results to a $\frac{1}{2}$-approximate Nash equilibrium for any 2-player game. Daskalakis *et al.* [21] also looked at the more demanding notion of *well supported approximate Nash equilibria* introduced in [20] and present an interesting reduction (of the same problem) to *win-lose* games (that is, games with all utilities equal to 0 and 1). For this more demanding notion, Daskalakis *et al.* showed that an approximation of $\frac{5}{6}$ is possible contingent upon a graph-theoretic conjecture.

Chen *et al.* [13] establish strong inapproximability results for approximate Nash equilibria. Their results imply that it is unlikely to obtain a fully polynomial time approximation scheme for Nash equilibria (unless $\mathcal{PPAD} \subseteq \mathcal{P}$).

1.6.4 Correlated Equilibria

Nash equilibrium [73, 74] is widely accepted as the standard notion of rationality in Game Theory. However, there are several other competing formulations of rationality; chief among them is the *correlated equilibrium,* proposed by Aumann [7]. Observe that the mixed Nash equilibrium is a distribution on the strategy space that is *uncorrelated* or *independent*; that is, it is the product of independent probability distributions, one for each player. In sharp contrast, a *correlated equilibrium* is a *general* distribution over strategy profiles. It must, however, possess an equilibrium property: If a strategy profile is drawn according to this distribution, and each player is told separately his suggested strategy (that is, his own component in the profile), then no player has an incentive to switch to a different strategy (assuming that all other players also obey), because the suggested strategy is the best in expectation. Correlated equilibria enjoy a very nice combinatorial structure: the set of correlated equilibria of a multi-player, non-cooperative game is a convex polytope, and all Nash equlibria are not only included in this polytope but they all lie on the boundary of the polytope. (See [75] for an elegant elementary proof of this latter result.)

As noted in the own words of Papadimitriou in his paper [82], the correlated equilibrium has several important advantages: It is a perfectly reasonable, simple and plausible concept; it is guaranteed to always exist (simply because the Nash equilibrium is a particular case of a correlated equilibrium); and it can be found in polynomial time for any number of players and strategies by linear programming, since the inequalities specifying the satisfaction of all players are linear. In fact, it turns out that the correlated equilibrium that optimizes any linear function of the players' utilities (for example, their sum) can be computed in polynomial time.

Succinct Games. Equilibria in games, of which the correlated equilibrium is a prominent example, are objects worth of studying from the algorithmic point of view. *Multiplayer games* are the most compelling specimens in this regard. But, to be of algorithmic interest, they must be *represented succinctly*. Succinct representation is required since otherwise a typical (multiplayer) game would need an exponential size of bits in order to be described. Some well known games that admit a succinct representation include:

- Symmetric games, where all players are identical and indistinguishable,
- *Graphical* games [51], where the players are the vertices of a graph, and the payoff for each player only depends on its own strategy and those of its neighbours;
- Congestion games, where the payoff of each player only depends on its strategy and those choosing the same strategy as him.

Papadimitriou and Roughgarden [83] initiated the systematic study of algorithmic issues involved in finding equilibria (both Nash and correlated) in games with a large number of players, which are succinctly represented. The authors develop a general framework for obtaining polynomial-time algorithms for optimizing over correlated equilibria in such settings. They show how such algorithms can be applied successfully to symmetric games, graphical games and congestion games, among others. They also present complexity results, implying that such algorithms are not in sight for certain other similar games. Finally, a polynomial-time algorithm, based on *quantifier elimination*, for finding a Nash equilibrium in symmetric games (when the number of strategies is relatively small) was presented.

Daskalakis and Papadimitriou [22] studied from the complexity point of view the problem of finding equilibria in games played on highly regular graphs with extremely succinct representation, such as the *d-dimensional grid*. There, it is argued that such games are of interest in modeling large systems of interacting agents. It has been shown by Daskalakis and Papadimitriou [22] that the problem of determining whether such a game on the d-dimensional grid has a pure Nash equilibrium depends on d, and the dichotomy is remarkably sharp: It is polynomial time solvable when $d = 1$, but $\mathcal{NEXP}$-complete for $d \geq 2$. In contrast, it was also proved that mixed Nash equilibria can be found in deterministic exponential time for any fixed d by quantifier elimination.

Recently, Papadimitriou [82] considered, and largely settled, the question of the existence of polynomial-time algorithms for computing correlated equilibria in succinctly representable multiplayer games. Papadimitriou developed a polynomial-time algorithm for finding correlated equilibria in a broad class of succinctly representable multiplayer games, encompassing essentially all kinds of such games we mentioned before.

The algorithm presented by Papadimitriou [82] was based on a careful mimicking of the existence proof due to Hart and Schmeidler [47], combined with an argument based on linear programming duality and the ellipsoid algorithm, Markov chain steady state computations, as well as application-specific methods for computing multivariate expectations.

1.7 DISCUSSION

In this chapter, we attempted a glimpse at the fascinating field of *Algorithmic Game Theory*. This is a field that is currently undergoing a very intense investigation by the community of the *Theory of Computing*. Although some fundamental theoretical questions have been resolved (for example, the complexity of computing Nash equilibria for 2-player games), there are still a lot of challenges ahead of us. Among those, most important are, in our opinion, the further complexity classification of algorithmic problems in Game Theory, and the further application of systematic techniques from Game Theory to modeling and evaluating modern computer systems with selfish entities.

REFERENCES

1. H. Ackermann, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*. IEEE Press, 2006. to appear.
2. H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash Equilibria in Player-Specific and Weighted Congestion Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE 2006)*, Lecture Notes in Computer Science. Springer, 2006. to appear.
3. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS 2006)*, volume 3884 of *Lecture Notes in Computer Science*, pages 218–229. Springer, 2006.
4. I. Althofer. On Sparse Approximations to Randomized Strategies and Convex Combinations. *Linear Algebra and its Applications*, 199:339–355, 1994.
5. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004)*, pages 295–304. IEEE Press, 2004.
6. J. Aspnes, K. Chang, and A. Yampolskiy. Inoculation Strategies for Victims of Viruses and the Sum-of-squares Partition Problem. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 43–52. Society for Industrial and Applied Mathematics, 2005.
7. R. J. Aumann. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
8. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 57–66. ACM Press, 2005.

9. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. *Theoretical Computer Science*, 361(2–3):200–209, 2006.
10. I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Taxes for Linear Atomic Congestion Games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA 2006)*, 2006. to appear.
11. X. Chen and X. Deng. 3-Nash is $\mathcal{PPAD}$-Complete. Technical Report TR05-134, Electronic Colloquium in Computational Complexity (ECCC), 2005.
12. X. Chen and X. Deng. Settling the Complexity of 2-Player Nash-Equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Press, 2006. to appear.
13. X. Chen, X. Deng, and S. Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, 2006. to appear.
14. E. R. Cheswick and S. M. Bellovin. *Firewalls and Internet Security*. Addison-Wesley, 1994.
15. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 67–73. ACM Press, 2005.
16. R. Cole, Y. Dodis, and T. Roughgarden. Pricing Network Edges for Heterogeneous Selfish Users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC 2003)*, pages 521–530. ACM Press, 2003.
17. A. Czumaj, P. Krysta, and B. Vöcking. Selfish Traffic Allocation for Server Farms. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC 2002)*, pages 287–296. ACM Press, 2002.
18. A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2002)*, pages 413–420. Society for Industrial and Applied Mathematics, 2002.
19. C. Daskalakis, A. Fabrikant, and C. H. Papadimitriou. The Game World Is Flat: The Complexity of Nash Equilibria in Succinct Games. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, volume 4051 of *Lecture Notes in Computer Science*, pages 513–524. Springer, 2006.
20. C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006)*, pages 71–78. ACM Press, 2006.
21. C. Daskalakis, A. Mehta, and C. Papadimitriou. A Note on Approximate Nash Equilibria. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE 2006)*, Lecture Notes in Computer Science. Springer, 2006. to appear.
22. C. Daskalakis and C. H. Papadimitriou. The Complexity of Equilibria in Highly Regular Graph Games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA 2005)*, volume 3669 of *Lecture Notes in Computer Science*, pages 71–82. Springer, 2005.
23. C. Daskalakis and C. H. Papadimitriou. Three-Player Games are Hard. Technical Report TR05-139, Electronic Colloquium in Computational Complexity (ECCC), 2005.
24. D. Dumrauf and M. Gairing. Price of Anarchy for Polynomial Wardrop Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE 2006)*, Lecture Notes in Computer Science. Springer, 2006. to appear.
25. J. Egerváry. Matrixok Kombinatorius Tulajdonságairól. *Matematikai és Fizikai Lapok*, 38:16–28, 1931.
26. R. Elsässer, M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. A Simple Graph-Theoretic Model for Selfish Restricted Scheduling. In *Proceedings of the 1st International Workshop on*

Internet and Network Economics (WINE 2005), volume 3828 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 2005.

27. A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004)*, pages 604–612. ACM Press, 2004.

28. J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the Cost of Muliticast Transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

29. S. Fischer and B. Vöcking. Adaptive Routing with Stale Information. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC 2005)*, pages 276–283. ACM Press, 2005.

30. S. Fischer and B. Vöcking. On the Structure and Complexity of Worst-case Equilibria. In *Proceedings of the 1st Workshop on Internet and Network Economics (WINE 2005)*, volume 3828 of *Lecture Notes in Computer Science*, pages 151–160. Springer Verlag, 2005.

31. L. Fleischer, K. Jain, and M. Mahdian. Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004)*, pages 277–285. IEEE Press, 2004.

32. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, 348(2–3):226–239, 2005.

33. D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms (WAOA 2005)*, volume 3879 of *Lecture Notes in Computer Science*, pages 161–175. Springer, 2006.

34. D. Fotakis, S. C. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP 2002)*, volume 2380 of *Lecture Notes in Computer Science*, pages 123–134. Springer, 2002.

35. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004)*, pages 613–622. ACM Press, 2004.

36. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Polynomial Social Cost. In *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS 2004)*, volume 3153 of *Lecture Notes in Computer Science*, pages 574–585. Springer, 2004.

37. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP 2004)*, volume 3142 of *Lecture Notes in Computer Science*, pages 645–657. Springer, 2004.

38. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. G. Spirakis. Extreme Nash Equilibria. In *Proceedings of the 8th Italian Conference of Theoretical Computer Science (ICTCS 2003)*, volume 2841 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2003.

39. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. G. Spirakis. Structure and Complexity of Extreme Nash Equilibria. *Theoretical Computer Science*, 343(1–2):133–157, 2005. Special Issue titled "Game Theory Meets Theoretical Computer Science", M. Mavronicolas and S. Abramsky guest eds.

40. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. In *Proceedings of the 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures*

(SPAA 2005), pages 203–212. ACM Press, 2005. Extended version accepted to *Theory of Computing Systems*, Special Issue with selected papers from the *17th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2005).*

41. *M. Gairing, B. Monien, and K. Tiemann. Routing (Un-)Splittable Flow in Games with Player-Specific Linear Latency Functions. In* Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, volume 4051 of* Lecture Notes in Computer Science*, pages 501–512. Springer, 2006.*
42. *J. Geanakoplos. Nash and Walras Equilibrium via Brouwer.* Economic Theory*, 2(2–3):585–603, 2003.*
43. *C. Georgiou, T. Pavlides, and A. Philippou. Network Uncertainty in Selfish routing. In* CD-ROM Proceedings of the 20th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2006)*, 2006.*
44. *P. W. Goldberg and C. H. Papadimitriou. Reducibility Among Equilibrium Problems. In* Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006)*, pages 61–70. ACM Press, 2006.*
45. *G. Gottlob, G. Greco, and F. Scarcello. Pure Nash Equilibria: Hard and Easy Games.* Journal of Artificial Intelligence Research*, 24:357–406, 2005.*
46. *J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III.* Management Science*, 14:159–182, 320–332, 468–502, 1967.*
47. *S. Hart and D. Schmeidler. Existence of Correlated Equilibria.* Mathematics of Operations Research*, 14(1):18–25, 1989.*
48. *M. Hirsch, C. H. Papadimitriou, and S. Vavasis. Exponential Lower Bounds for Finding Brouwer Fixpoints.* Journal of Complexity*, 5:379–41, 1989.*
49. *D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How Easy is Local Search?* Journal of Computer and System Sciences*, 17(1):79–100, 1988.*
50. *A. Kaporis and P. Spirakis. The Price of Optimum in Stackelberg Games. In* Proceedings of the 18th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2006)*, pages pp. 19–28, 2006.*
51. *M. Kearns, M. Littman, and S. Singh. Graphical Models for Game Theory. In* Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, page 253–260, 2001.*
52. *M. Kearns and L. Ortiz. Algorithms for Interdependent Security Games. In* Proceedings of the 16th Annual Conference on Neural Information Processing Systems (NIPS 2004)*, pages 288–297. MIT Press, 2004.*
53. *J. Kleinberg and É. Tardos.* Algorithm Design*. Addison-Wesley, 2005.*
54. *D. König. Graphok és Matrixok.* Matematikai és Fizikai Lapok*, 38:116–119, 1931.*
55. *S. Kontogiannis and P. Spirakis. Atomic Selfish Routing in Networks: a Survey. In* Proceedings of the 1st International Workshop on Internet and Network Economics (WINE 2005)*, volume 3828 of* Lecture Notes in Computer Science*, pages 989–1002. Springer, 2005.*
56. *S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. In* Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE 2006)*, Lecture Notes in Computer Science. Springer, 2006. to appear.*
57. *Y. A. Korilis, A. Lazar, and A. Orda. The Designer's Perspective to Noncooperative Networks. In* Proceedings of the 14th Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE INFOCOM 1995)*, volume 2, pages 562–570, 1995.*
58. *Y. A. Korilis, A. Lazar, and A. Orda. Achieving Network Optima using Stackelberg Routing Strategies.* IEEE/ACM Transactions of Networking*, 5(1):161–173, 1997.*

59. *E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. Approximate Equilibria and Ball Fusion.* Theory of Computing Systems, *36(6):683–693, 2003.*

60. *E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In* Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS 1999), *volume 1563 of* Lecture Notes in Computer Science, *pages 404–413. Springer, 1999.*

61. *C.E. Lemke and Jr. J.T. Howson. Equilibrium Points of Bimatrix Games.* Journal of the Society for Industrial and Applied Mathematics, *12:413–423, 1964.*

62. *R. J. Lipton, E. Markakis, and A. Mehta. Playing Large Games Using Simple Strategies. In* Proceedings 4th ACM Conference on Electronic Commerce (EC-2003), *pages 36–41. ACM Press, 2003.*

63. *T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In* Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS 2004), *volume 2996 of* Lecture Notes in Computer Science, *pages 547–558. Springer, 2004.*

64. *M. Mavronicolas, L. Michael, V. G. Papadopoulou, A. Philippou, and P. G. Spirakis. The Price of Defense. In* Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science (MFCS 2006), *volume 4162 of* Lecture Notes in Computer Science, *pages 717–728. Springer, 2006.*

65. *M. Mavronicolas, P. Panagopoulou, and P. Spirakis. A Cost Mechanism for Fair Pricing of Resource Usage. In* Proceedings of the 1st International Workshop on Internet and Network Economics (WINE 2005), *volume 3828 of* Lecture Notes in Computer Science, *pages 210–224. Springer, 2005.*

66. *M. Mavronicolas, V. Papadopoulou, G. Persiano, A. Philippou, and P. Spirakis. The Price of Defense and Fractional Matchings. In* Proceedings of the 8th International Conference on Distributed Computing and Networking (ICDCN 2006), *2006. to appear.*

67. *M. Mavronicolas, V. G. Papadopoulou, A. Philippou, and P. G. Spirakis. A Graph-Theoretic Network Security Game. In* Proceedings of the 1st International Workshop on Internet and Network Economics (WINE 2005), *volume 3828 of* Lecture Notes in Computer Science, *pages 969–978. Springer, 2005.*

68. *M. Mavronicolas, V. G. Papadopoulou, A. Philippou, and P. G. Spirakis. A Network Game with Attacker and Protector Entities. In* Proceedings of the 16th Annual International Symposium on Algorithms and Computation (ISAAC 2005), *volume 3827 of* Lecture Notes in Computer Science, *pages 288–297. Springer, 2005.*

69. *M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In* Proceedings of the 33th Annual ACM Symposium on Theory of Computing (STOC 2001), *pages 510–519. ACM Press, 2001. Full version accepted to Algorithmica.*

70. *I. Milchtaich. Congestion Games with Player-Specific Payoff Functions.* Games and Economic Behavior, *13(1):111–124, 1996.*

71. *D. Monderer and L. S. Shapley. Potential Games.* Games and Economic Behavior, *14(1):124–143, 1996.*

72. *T. Moscibroda, S. Schmid, and R. Wattenhofer. When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In* Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC 2006). *ACM Press, 2006.*

73. *J. F. Nash. Equilibrium Points in N-Person Games. In* Proceedings of National Acanemy of Sciences of the United States of America, *volume 36, pages 48–49, 1950.*

74. *J. F. Nash. Non-cooperative Games.* Annals of Mathematics, *54(2):286–295, 1951.*

75. *R. Nau, S. G. Canovas, and P. Hansen. On the Geometry of Nash Equilibria and Correlated Equilibria.* International Journal of Game Theory, *32(4):443–453, 2003.*

76. *N. Nisan and A. Ronen. Algorithmic Mechanism Design.* Games and Economic Behavior, *35(1-2):166–196, 2001.*

77. *M. Osborne.* An Introduction to Game Theory. *Oxford University Press, 2003.*

78. *M. Osborne and A. Rubinstein.* A Cousre in Game Theory. *MIT Press, 1994.*

79. *P. Panagopoulou and P. Spirakis. Efficient Convergence to Pure Nash Equilibria in Weighted Network Congestion Games. In* Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms (WEA 2005)*, volume 3503 of* Lecture Notes in Computer Science*, pages 203–215. Springer, 2005.*

80. *C. H. Papadimitriou.* Computational Complexity. *Addison-Wesley, 1994.*

81. *C. H. Papadimitriou. Algorithms, Games, and the Internet. In* Proceedings of the 33th Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 749–753. ACM Press, 2001.*

82. *C. H. Papadimitriou. Computing Correlated Equilibria in Multi-Player Games. In* Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 49–56. ACM Press, 2005.*

83. *C. H. Papadimitriou and T. Roughgarden. Computing Equilibria in Multi-Player Games. In* Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 82–91. Society for Industrial and Applied Mathematics, 2005.*

84. *R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria.* International Journal of Game Theory*, 2(65–67), 1973.*

85. *T. Roughgarden. Stackelberg Scheduling Strategies.* SIAM Journal of Computing*, 33(2):332–350, 2003.*

86. *T. Roughgarden.* Selfish Routing and the Price of Anarchy. *The MIT Press, 2005.*

87. *T. Roughgarden. On the Severity of Braess's Paradox: Designing Networks for Selfish Users is Hard.* Journal of Computer and System Sciences*, pages 922–953, 2006.*

88. *T. Roughgarden and É. Tardos. How Bad is Selfish Routing?* Journal of the ACM*, 49(2):236–259, 2002.*

89. *E. H. Spafford. The Internet Worm: Crisis and Aftermath.* Communications of the ACM, *6(2–3):678–687, 1989.*

90. *W. Stallings.* Cryptography and Network Security: Principles and Practice. *Prentice Hall, Third edition, 2003.*