Chapter 9
Game-Theoretic Approaches to Optimization Problems in Communication Networks

Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, Michele Flammini, Christos Kaklamanis, Gianpiero Monaco, and Luca Moscardelli

Abstract In this chapter we consider fundamental optimization problems arising in communication networks. We consider scenarios where there is no central authority that coordinates the network users in order to achieve efficient solutions. Instead, the users act in an uncoordinated and selfish manner and reach solutions to the above problems that are consistent only with their selfishness. In this sense, the users act aiming to optimize their own objectives with no regard to the globally optimum system performance. Such a behavior poses several intriguing questions ranging from the definition of reasonable and practical models for studying it to the quantification of the efficiency loss due to the lack of users' cooperation. We present several results we achieved recently in this research area and propose interesting future research directions.

Key words: non-cooperative networks, strategic games, Nash equilibria, price of anarchy, price of stability

---

Vittorio Bilò
Department of Mathematics, University of Salento, Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce, Italy, e-mail: `vittorio.bilo@unile.it`

Ioannis Caragiannis · Christos Kaklamanis
Computer Technology Institute & Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, e-mail: `{caragian,kakl}@ceid.upatras.gr`

Angelo Fanelli · Michele Flammini · Gianpiero Monaco · Luca Moscardelli
Dipartimento di Informatica, Università degli Studi dell'Aquila, Via Vetoio, Loc. Coppito, 67100 L'Aquila, Italy, e-mail: `{angelo.fanelli,flammini,gianpiero.monaco,moscardelli}@di.univaq.it`

9.1 Introduction

In the last few years, a mutual interest between the community of computer scientists and of economists is emerging, especially concerning the relationships between Distributed Systems and Game Theory. The trigger for this synergy is the evolution of uncoordinated and non-cooperative communication networks such as the Internet, peer-to-peer networks, and wireless ad hoc networks, and the consequent diversity of the owners of the relative physical components (routers, computing devices, communication channels, etc.). This has determined a radical change in the definition of what distributed computing means. Indeed, in the classical computational model adopted in Theoretical Computer Science, processors are faithful executors of an algorithm. Conversely, in the incentive-based model, which is the one traditionally studied by economists, processors (usually called agents) pursue a selfish strategy and the system evolves as a consequence of the complex interactions among them. Therefore, Game Theory comes to pursue the synthesis between two requirements: the system designer's or operator's (whose goal is to compute a socially efficient solution), and the agents' (which, in a non-cooperative scenario, aim to maximize their own profit and, therefore, could induce the system to reach suboptimal solutions).

Let us introduce informally the concepts studied in this chapter with an example. Assume that we have a simple communication network in which two nodes s and t are connected through two directed links e_1 and e_2 (see Figure 9.1). There are two users, each having a unit amount of traffic. Each user wishes to route her traffic from s to t unsplit, using one of the two links. Link e_1 has a latency function $d_{e_1}(x) = x$, which means that the latency experienced by each user that uses this link is x when x users in total use this link. The latency function of link e_2 is $d_{e_2}(x) = 2+\varepsilon$, where ε is a negligible but strictly positive constant. A natural objective is to minimize the average (or total) latency among the two users. So, from the network designer point of view, a solution which forces one user to use link e_1 and the other to use link e_2 would be the desired one with a total latency of $3+\varepsilon$. However, if each user is selfish in the sense that she prefers to use the link which minimizes her latency given the decision of the other user, we could never end up with this desired solution. The user that is assigned to link e_2 (where she experiences a latency of $2+\varepsilon$) has an incentive to change her strategy and use link e_1 instead; then, the latency experienced would be only 2. From this latter assignment, no user has an incentive to unilaterally deviate; such a solution is consistent with the selfish nature of the behavior of the users. Unfortunately, selfishness may result in deterioration of system performance; the total latency in this latter solution is 4 since both users experience a latency of 2 on link e_1.

In the above example, we have implicitly assumed that users (called players in the following) play a strategic game which reflects their selfish behavior. Solutions that are consistent with this behavior are called equilibria; these notions are formally defined in the next section. Issues related to the existence, the structure, the computation, and the efficiency of equilibria in strategic games arising from optimization problems are studied (among other issues) in the recently emerging field

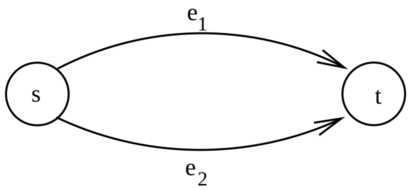

Fig. 9.1 A simple communication network

of Algorithmic Game Theory. In this chapter we present some basic contributions in this fascinating emerging area related to some fundamental optimization problems that arise in communication networks with non-cooperative users. We focus on results which we have obtained recently; the interested reader may see [39] for a more systematic study of the recent trends in Algorithmic Game Theory.

9.2 Preliminary Notions

We begin with some necessary definitions. Given a set U, a k-tuple $a = (u_1, \dots, u_k)$ of elements of U, an index $i \in \{1, \dots, k\}$, and an element $u \in U$, we write $(a_{-i}, u) = (u_1, \dots, u_{i-1}, u, u_{i+1}, \dots, u_k)$ to denote the k-tuple obtained from a by replacing u_i with u.

Strategic games. A strategic game $\mathcal{G}$ is a triple $\mathcal{G} = (P, S_{i\in P}, \omega_{i\in P})$ where P is a set of n players, S_i is the set of strategies available to player i and $\omega_i : S_1 \times \dots \times S_n \mapsto \mathrm{R}$ is her payoff function. The payoff function ω_i can model either a benefit or a cost for player i; thus each player may seek the maximization or the minimization of her payoff. In the sequel we will always assume that the payoff functions model costs for the players.

States and improving steps. The set $S = S_1 \times \dots \times S_n$ is called the set of states or strategy profiles of $\mathcal{G}$. Consider a state $\sigma = (\sigma_1, \dots, \sigma_n) \in S$. A player i cannot be happy with her payoff in σ if there exists another state $\sigma' = (\sigma_{-i}, s)$, for some $s \in S_i$, such that $\omega_i(\sigma') < \omega_i(\sigma)$. The action of changing their strategy from σ_i to s (with the consequent transition of $\mathcal{G}$ from state σ to state σ') is called an improving step or a selfish move performed by player i.

Pure Nash equilibria. A very important and challenging issue in game theory is to characterize solutions of games which are consistent with the selfish and rational behavior of the players. Among all the approaches proposed so far, the notion of Nash equilibrium is the most accepted and used one. A state σ is a pure Nash equilibrium if no player has an improving step, that is, $\forall i \in P$ and $\forall s \in S_i$ it holds $\omega_i(\sigma_{-i}, s) \geq \omega_i(\sigma)$. A pure Nash equilibrium is a stable outcome of a game in the sense that it is a state in which all players are satisfied with their payoff and none of them can improve it by unilaterally changing her strategy. The drawback of this notion is, however, due to the fact that pure Nash equilibria are not guaranteed to exist for any game.

Mixed strategies. A mixed strategy for player i is a probability distribution Y_i defined over the set S_i of her pure strategies. Given a mixed strategy Y_i for player i, the support of Y_i, denoted as $support(Y_i)$, is the set of strategies $s \in S_i$ for which $Y_i(s) > 0$. A mixed strategy $Y = (Y_1 \ldots, Y_n)$ is an n-tuple of mixed strategies Y_i for each player $i \in P$. The support of a mixed strategy Y is defined as the set of states $support(Y) = support(Y_1) \times \ldots \times support(Y_n)$. The payoff of player i yielded by a mixed strategy Y is defined as $\omega_i(Y) = \sum_{\sigma \in support(Y)} (\omega_i(\sigma) \cdot Prob_Y(\sigma))$, where $Prob_Y(\sigma) = \prod_{i=1}^{n} Y_i(\sigma_i)$.

Mixed Nash equilibria. A mixed strategy Y is a (mixed) Nash equilibrium if $\forall i \in P$ and for any probability distribution Z defined over S_i it holds $\omega_i(Y_{-i}, Z) \geq \omega_i(Y)$. A fully mixed Nash equilibrium Y is a mixed Nash equilibrium such that $support(Y_i) = S_i$ for each $i \in P$. Mixed Nash equilibria are, hence, a generalization of pure ones, obtained by allowing players to randomize on their chosen strategies and computing the resulting payoffs in expectation. In this case the solution concept becomes less reasonable (since, in practice, we are always asked to choose pure strategies); but, on the other hand, Nash's famous theorem [38] guarantees the existence of at least one mixed Nash equilibrium for any game.

Price of anarchy and stability. One of the main concerns when dealing with non-cooperative systems is to bound their inefficiencies due to the lack of coordination among the players. More formally, given a function $\gamma : S \mapsto \mathbf{R}$, called the social function, let σ^* be a state optimizing γ. On the other hand, given a mixed strategy Y, the social value of Y is defined as $\gamma(Y) = \sum_{\sigma \in support(Y)} (\gamma(\sigma) \cdot Prob_Y(\sigma))$. The price of anarchy [35] of game $\mathscr{G}$ for the social function γ is defined as $PoA(\mathscr{G}, \gamma) = \sup_{Y \in NE(\mathscr{G})} \frac{\gamma(Y)}{\gamma(\sigma^*)}$, while the price of stability [3] of $\mathscr{G}$ for the social function γ is defined as $PoS(\mathscr{G}, \gamma) = \inf_{Y \in NE(\mathscr{G})} \frac{\gamma(Y)}{\gamma(\sigma^*)}$, where $NE(\mathscr{G})$ denotes the set of Nash equilibria of $\mathscr{G}$. By restricting $NE(\mathscr{G})$ to the set of pure Nash equilibria, similar definitions for the price of anarchy and stability are obtained for pure strategies. The price of anarchy is a classical worst-case analysis and it measures the loss of performance due to the selfish behavior of players in this case. On the other hand, the price of stability gives us information on the minimum loss of performance a non-cooperative system has to suffer. This issue is best understood when dealing with games, such as network design games, in which it is possible to assume the presence of a central authority proposing or imposing a pure Nash equilibrium to the players. In this setting, the best Nash equilibrium represents the optimal solution among all the proposed ones from which no player would have an incentive to defect. By extending the above definition, the price of anarchy after a best response walk can be defined as the worst-case ratio of the social value of the last state of the walk to the value of the social optimum.

Nash dynamics graph. For any game $\mathscr{G}$ it is possible to define a graph $D(\mathscr{G}) = (S,A)$ capturing the dynamic behavior of the players in the game. The set of nodes of $D(\mathscr{G})$ is represented by the set of states of $\mathscr{G}$, and there is a directed edge between any two nodes σ and σ' with label i if and only if player i has an improving step from σ to σ' in $\mathscr{G}$. By definition, we have that $\mathscr{G}$ possesses pure Nash equilibria if and only if $D(\mathscr{G})$ has sink nodes, that is, nodes without any outgoing edge. An important issue related to the notion of pure Nash equilibria is the finite improvement path (FIP) property. A game has this property if, starting from any initial state and letting a player perform arbitrary improving steps, a pure Nash equilibrium is always reached, i.e., there does not exist an infinite sequence of improving steps. Again, by definition, we have that $\mathscr{G}$ has the FIP property if and only if $D(\mathscr{G})$ is acyclic, that is, a DAG (directed acyclic graph). In a given state a player i may have more than one improving step, i.e., for a given node there may exist more than one outgoing edge labeled i in $D(\mathscr{G})$. An improving step for player i in state σ is a best response for i in σ if it yields the maximum improvement in i's payoff among all improving steps i has in σ. We define the best response dynamics graph of $\mathscr{G}$ as the graph obtained from $D(\mathscr{G})$ by removing all edges not modeling best responses.

Best response walks. Any path in the best response dynamics graph is called a *best response walk*. As there are games that do not have the FIP property as well as games for which the number of best responses (or, more generally, improving steps) needed to reach an equilibrium starting from an arbitrary state may be exponentially large, it becomes important to evaluate the social value of states reached after a finite number of selfish moves. To this aim, Mirrokni and Vetta [36] introduced the following models capturing the intuitive notion of a fair sequence of moves.

Covering walk. A walk in the best response dynamics graph is a covering walk if for each player i, it has at least one edge with label i.
k-Covering walk. A walk in the best response dynamics graph is a k-covering walk if it can be split into k disjoint covering walks.
One-round walk. A covering walk of length n in the best response dynamics graph is a one-round walk, that is, each player performs exactly one best response.

Potential games. If $\mathscr{G}$ has the FIP property then, by exploiting the topological ordering of the nodes in $D(\mathscr{G})$, it is possible to define a potential function for $\mathscr{G}$, that is, a function $\Phi : S \mapsto \mathbb{R}$ always decreasing or always increasing each time an improving step is performed. By generalizing this observation, it is possible to define the class of potential games as the class of all games having a potential function and, hence, the FIP property. Depending on the properties satisfied by Φ, three different kinds of potential games can be defined.

1. *Exact potential* games, where $\forall \sigma \in S$ and $\forall s \in S_i$, it holds $\Phi(\sigma) - \Phi(\sigma_{-i}, s) = \omega_i(\sigma) - \omega_i(\sigma_{-i}, s)$.

2. *Weighted potential* games, where $\exists\beta = (\beta_1,\ldots,\beta_n)$ such that $\forall\sigma \in S$ and $\forall s \in S_i$, it holds $\Phi(\sigma) - \Phi(\sigma_{-i},s) = \beta_i(\omega_i(\sigma) - \omega_i(\sigma_{-i},s))$.
3. *Ordinal potential* games, where $\forall\sigma \in S$ and $\forall s \in S_i$, it holds that $\Phi(\sigma) - \Phi(\sigma_{-i},s)$ and $\omega(\sigma) - \omega(\sigma_{-i},s)$ have the same sign.

Congestion games. Perhaps the most famous class of games is that of congestion games. Results concerning congestion games are presented in Sections 9.3 and 9.7. Here, we present the basic related definitions. A congestion game is a 4-tuple $(P,E,S_{i\in P},d_{e\in E})$, where P is a set of n players, E is a set of m resources, $S_i \subseteq 2^{|E|}$ for each $i \in P$, and $d_e : \mathrm{N} \mapsto \mathbf{R}$ for each $e \in E$. Strictly speaking, each player in a congestion game can choose among different subsets of resources; each resource e has an associated delay (latency) function d_e returning the delay experienced by any player using e in terms of the number of players using it. Once $n_e(\sigma) = |\{i \in P : e \in \sigma_i\}|$ is defined as the number of players using resource e in state σ, the payoff function of player i is defined as $\omega_i(\sigma) = \sum_{e\in\sigma_i} d_e(n_e(\sigma))$. Congestion games were introduced by Rosenthal [41], who proved that they have the FIP property by defining the potential function $\Phi(\sigma) = \sum_{e\in E}\sum_{i=1}^{n_e(\sigma)} d_e(i)$. As shown by Monderer and Shapley [37], the class of congestion games is equivalent to that of exact potential games. The class of weighted congestion games is a generalization obtained by allowing players to have different demands. In this case the congestion of each resource e is defined as $n_e(\sigma) = \sum_{i\in P:e\in\sigma_i} r_i$, where r_i is the demand of player i. Some special cases of congestion games are particularly interesting to computer scientists. In network congestion games, resources correspond to the links of a network. Each player i has a source s_i and the destination t_i and her set of strategies corresponds to the set of paths connecting s_i to t_i. Load balancing games are congestion games in which all possible strategies for each player are singletons.

Cost sharing games. Similarly to congestion games, a cost sharing game is a 4-tuple $(P,E,S_{i\in P},c_{e\in E})$, where P is a set of n players, E is a set of m resources, $S_i \subseteq 2^{|E|}$ for each $i \in P$, and $c : E \mapsto \mathbf{R}^+$. In these games, each resource e has a certain fixed cost c_e, which has to be shared among the players. In particular, each strategy profile σ requires the use of a set of resources $\Pi(\sigma) = \bigcup_{i=1}^{n}\{e \in E : e \in \sigma_i\}$ of total cost $cost(\Pi(\sigma)) = \sum_{e\in\Pi(\sigma)} c_e$, and the payoffs yielded by σ are computed by applying a certain cost sharing method $\mathscr{M}$ distributing $cost(\Pi(\sigma))$ among the players. Depending on the definition of $\mathscr{M}$, different games can be obtained. The most important family of cost sharing games arises in network design where the cost of a certain self-emerging network needs to be shared among its users. In this case, there is an edge-weighted graph $G = (V,E,c)$, with $c : E \mapsto \mathbf{R}^+$, modeling the set of links which can be potentially implemented, as well as their building costs. Each player i has a source s_i and a destination t_i that she wants to connect, and her set of strategies corresponds to the set of paths connecting s_i to t_i in G. This general network cost sharing game is referred to as the multi-commodity case. An interesting and much studied restriction is the single-commodity case, also called the multicast case, in which there is a common source s for all players and a set

$R \subset V \setminus \{s\}$ of n receivers which need to be connected to s. The most-used cost sharing method is the one based on the well-known Shapley value formula [45] which equally distributes the cost of each resource among the players using it, thus yielding $\omega_i(\sigma) = \sum_{e\in\sigma_i} \frac{c_e}{n_e(\sigma)}$. A cost sharing game in which the underlying cost sharing method is the Shapley value is called a Shapley cost sharing game. As first noted in [2] these games are also congestion games. Results related to cost sharing games are covered in Sections 9.4 and 9.7.

9.3 Congestion Games

In this section we consider congestion games. Besides general congestion games, we also focus on network congestion games and load balancing games.

We study issues related to the efficiency of equilibria in congestion games according to the social cost function of the weighted total latency, i.e., the sum of the latency experienced by each player multiplied by her demand. Formally, the social cost of a state σ is $\gamma(\sigma) = \sum_i r_i \omega_i(\sigma)$. An equivalent definition is $\gamma(\sigma) = \sum_e n_e(\sigma) d_e(n_e(\sigma))$, since the sum of the demands of the players using resource e is $n_e(\sigma)$ and each of them experiences a latency of $d_e(n_e(\sigma))$ on e.

Price of anarchy. A large part of the literature related to congestion games is devoted to providing upper and lower bounds on the price of anarchy. What we essentially want to do in this case is to relate the maximum social cost over all equilibria to the optimal social cost. The following simple argument is the heart of most related proofs in the literature. Consider an equilibrium σ for a congestion game and let σ^* be the assignment of optimal social cost. Assume that the latency functions are linear with the simple form $d_e(x) = \alpha_e x$ and that the demands of the players are the same (and equal to 1). Since no player has an incentive to change her strategy in σ, this also means that $\omega_i(\sigma) \leq \omega_i(\sigma_{-i}, \sigma_i^*)$, i.e., that the player i has no incentive to change her strategy to the one it uses in σ^*. Substituting the payoff, this yields that

$$\sum_{e\in\sigma_i} d_e(n_e(\sigma)) \leq \sum_{e\in\sigma_i^*} d_e(n_e(\sigma_{-i}, \sigma_i^*)) \leq \sum_{e\in\sigma_i^*} d_e(n_e(\sigma)+1).$$

The second inequality follows since $(\sigma_{-i}, \sigma_i^*)$ and σ differ only in the strategy of player i and since the latency functions are non-decreasing. Then, the social cost of σ is

$$\begin{aligned}\gamma(\sigma) &= \sum_i \omega_i(\sigma)\\ &= \sum_i \sum_{e\in\sigma_i} d_e(n_e(\sigma))\\ &\le \sum_i \sum_{e\in\sigma_i^*} d_e(n_e(\sigma)+1)\\ &= \sum_e \sum_{i:e\in\sigma_i^*} d_e(n_e(\sigma)+1)\\ &= \sum_e n_e(\sigma^*)d_e(n_e(\sigma)+1)\\ &= \sum_e \alpha_e\left(n_e(\sigma^*)n_e(\sigma)+n_e(\sigma^*)\right) \end{aligned} \tag{9.1}$$

The next step is to apply a simple inequality such as $xy+y \le \frac{1}{3}x^2+\frac{5}{3}y^2$ for any integer $x,y \ge 0$ on the right part of (9.1) to obtain (by setting $x = n_e(\sigma)$ and $y = n_e(\sigma^*)$)

$$\begin{aligned}\gamma(\sigma) &\le \sum_e \alpha_e\left(\frac{1}{3}n_e^2(\sigma)+\frac{5}{3}n_e^2(\sigma^*)\right)\\ &= \frac{1}{3}\gamma(\sigma)+\frac{5}{3}\gamma(\sigma^*)\end{aligned}$$

and, hence, $\gamma(\sigma) \le \frac{5}{2}\gamma(\sigma^*)$, which means that the price of anarchy is at most $5/2$.

The above bound was independently obtained by Christodoulou and Koutsoupias [21] and Awerbuch et al. [4]. The proof can be generalized to prove that the price of anarchy of weighted linear congestion games over mixed Nash equilibria is at most $\phi^2 \approx 2.62$, where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. These results are tight. Several papers in the literature (e.g., [1, 4, 21, 46]) have applied the same technique for more general (e.g., polynomial) latency functions.

Prior to these works, Suri et al. [46] had obtained the same bound for load balancing games. A natural question in load balancing games is whether their simplicity (compared to congestion games) leads to better bounds on the price of anarchy. In [15], it has been shown that this is not the case in general by proving tight lower bounds of 2.62 and $5/2$ on the price of anarchy of load balancing games with linear latency functions and players with different or equal demands, respectively. However, there are cases where load balancing games have some special structure that decreases the price of anarchy. For example, [15] exploits information about the structural properties of load balancing games on machines with identical linear latency functions to show an upper bound of 2.012 for the price of anarchy over pure Nash equilibria. This result is tight; a matching lower bound had been previously obtained in [46].

Price of stability. The argument that is used to prove upper bounds on the price of stability is similar. Besides the use of the Nash inequality, information provided by the potential function is also used. Essentially, we aim to relate the social cost of the best pure Nash equilibrium σ' to the social cost of the optimal assignment σ^*. We know that the social cost of σ' is upper-bounded by the social cost of an equilibrium σ which can be reached in the Nash dynamics graph by following improving steps from σ^*. Hence, in order to compute an upper bound on the price of stability, it suffices to compute an upper bound on the social cost of σ in terms of the optimal social cost of σ^*. By the definition of the potential function, we know that the potential of σ is smaller than the potential of σ^*, i.e., $\Phi_R(\sigma) \leq \Phi_R(\sigma^*)$. Again, let us assume that the latency functions have the simple linear form $d_e(x) = \alpha_e x$ and that the demands of the players are the same (and equal to 1). Then, the potential of σ can be written as

$$\Phi_R(\sigma) = \sum_e \sum_{j=1}^{n_e(\sigma)} d_e(j) = \frac{1}{2} \sum_e \alpha_e \left(n_e^2(\sigma) + n_e(\sigma)\right)$$

and the inequality of potentials can be used to obtain the following derivation.

$$\begin{aligned} \gamma(\sigma) &= \sum_e \alpha_e n_e^2(\sigma) \\ &\leq \sum_e \alpha_e \left(n_e^2(\sigma) + n_e(\sigma) - n_e(\sigma^*)\right) \end{aligned} \tag{9.2}$$

The main idea in the proofs of [22] and [15] is to use the information provided by both (9.1) and (9.2) in order to bound the social cost of σ in terms of the optimal social cost. This requires multiplying each of the two inequalities with a particular coefficient, summing them, and then using an inequality on integers in order to bound the right part by an expression that contains only the social cost of σ and the optimal social cost of σ^*. In this way, an upper bound of $1 + \frac{1}{\sqrt{3}} \approx 1.577$ on the price of stability is proved in [15] (prior to this result, a slightly inferior bound of 1.6 had been presented in [22]). This bound is tight; a congestion game with a single equilibrium of social cost $1 + \frac{1}{\sqrt{3}}$ times the optimal social cost has been presented in [22].

In load balancing games with linear latency functions, the price of stability has been proved to be much smaller, namely $4/3$. The lower bound easily follows by a simple game with two players of unit demands on two resources with latency functions $f_{e_1}(x) = x$ and $f_{e_2}(x) = (2+\varepsilon)x$ for arbitrarily small $\varepsilon > 0$. The state in which both players select resource e_1 is the only pure Nash equilibrium of social cost 4 while the states in which the players select different resources have social cost $3+\varepsilon$. In order to prove the upper bound, the information provided by the potential function is not enough. The approach followed in [15] is to provide a particular sequence of improving steps from a state σ^* of optimal social cost to a pure Nash equilibrium σ. The sequence is defined in such a way that a comparison between the social costs of σ and σ^* is feasible. The interested reader can inspect [15] for the details of the technique.

Independently, Anshelevich et al. [2] have shown a more general result that characterizes the price of stability of any network congestion game in which players have a common source (or a common destination). Namely, the result of [2] states that the price of stability of these games is at most the price of anarchy of corresponding nonatomic congestion games [44] on the same network. The main assumption in these games which differentiates them from the games we study in this section is that the number of players is infinite, and each is supposed to control a negligible amount of traffic from its source to its destination.

Performance of k-round walks. Convergence issues have been the subject of several papers. In [24], the authors show an upper bound of $\Theta(n)$ on the price of anarchy after one round, and a lower bound of $\sqrt[2^{\mathcal{O}(k)}]{n}/k$ after k rounds. Convergence to a constant price of anarchy in a polynomial number of random rounds can be inferred directly from the sink equilibria results of [36], while the $\sqrt[2^{\mathcal{O}(k)}]{n}/k$ lower bound of [24] implies that a number of rounds at least proportional to $\log\log n$ is necessary. In [27] it is shown that the price of anarchy achieved after k rounds is $\mathcal{O}(\sqrt[2^{k-1}]{n})$ while the lower bound of [24] is refined to $\Omega(\sqrt[2^{k-1}]{n}/k)$, which is asymptotically matching for constant values of k. As a consequence, $\log\log n$ rounds are not only necessary, but also sufficient in order to achieve a constant price of anarchy, i.e., comparable to the one at Nash equilibrium. [27] also provides a new lower bound of $\Omega(\sqrt[2^{k}-1]{n})$ for load balancing games, thus showing that a number of rounds proportional to $\log\log n$ is necessary and sufficient under such a restriction.

Coping with selfishness. Another line of research in congestion games aims to cope with selfishness. There are at least four different approaches that have been proposed in the literature: coordination mechanisms [6, 14, 23, 31], Stackelberg strategies [42], network design or resource removal [5, 43], and taxes or tolls [16, 17, 25, 29, 32]. The general idea behind each of them is to change the congestion game at hand in a reasonable way so that the resulting game has a small price of anarchy. In network design [43], some of the resources are removed; this also constrains the sets of strategies of the players. The results of Azar and Epstein [5] indicate that even deciding whether network design can decrease the price of anarchy of weighted network congestion games with linear latency functions to significantly less than ϕ^2 (the upper bound on the price of anarchy of such games) is NP-hard. Besides this, they show that there exist games in which network design is not helpful at all.

In the following, we briefly discuss the results in [16], which seems to be the first work dealing with taxes in atomic congestion games. In contrast, the study of taxes in nonatomic congestion games has a long history, starting at the beginning of the twentieth century with the work of Pigou [40]; and, besides recent contributions from Computer Science [25, 29, 32], it contains contributions from Economics and Transportation Science.

The model used in [16] is the following. A tax function $\delta : E \times Q^+ \to Q^+$ is introduced, meaning that a tax $\delta_e(w)$ is assigned to each player of demand w wishing to use the resource e. In this way, an extended game is obtained in which the cost of each player is the sum of the latency experienced and the tax she pays. The latency functions considered in [16] are linear. By defining a potential function (similar to the potential function for weighted linear congestion games [30]), it is shown that the extended game always has a pure Nash equilibrium. We note that network design can be thought of as a special kind of tax (infinite tax on some of the resources).

In the case of symmetric load balancing games, pure optimal taxes (i.e., taxes such that any pure Nash equilibrium of the extended game has optimal total latency) exist and can be constructed in polynomial time. Besides this positive result, even in simple congestion games, pure optimal taxes do not exist. Consider the following simple congestion game with four resources e_1, e_2, e_3, e_4 with the same latency function $d(x) = x$ and four players of demand 1: two long players, each having strategies $\{e_1, e_3\}, \{e_2, e_4\}$ and two short players, each having strategies $\{e_1\}, \{e_2\}$. Observe that assignments in which each of the resources e_1 and e_2 is used by one long and one short player have optimal total latency 10. Consider a tax assignment δ. It can be easily verified that if $\delta_{e_1} + \delta_{e_3} \leq \delta_{e_2} + \delta_{e_4}$, the assignment where the long players select strategy $\{e_1, e_3\}$ and the short players select strategy $\{e_2\}$ is a pure Nash equilibrium for the extended game; otherwise, the assignment where the long players select strategy $\{e_2, e_4\}$ and the short players select strategy $\{e_1\}$ is a pure Nash equilibrium for the extended game. In any case, the total latency is 12. Even simpler load balancing games do not admit pure optimal taxes; the reader can inspect [16] and [17] for more such negative statements.

Since taxes cannot always force optimal system performance, the next step is to study whether there are taxes such that the extended game has efficient equilibria. The efficiency of these equilibria (mixed or pure) is assessed in two different ways. In the case of refundable taxes, the social cost is simply the (weighted) total latency of the players. In the case of nonrefundable taxes, the social cost is the (weighted) total latency plus the total taxes paid. Assumptions similar to those in refundable taxes are common in the study of taxes for nonatomic congestion games in the Economics and Transportation literature and capture the scenarios where the collected taxes can be feasibly returned (directly or indirectly) to the players (e.g., as a "lump-sum refund").

The terms ρ-pure-efficient and ρ-mixed-efficient are used to refer to taxes which guarantee that the social cost of each equilibrium of the extended game is at most ρ times larger than the optimal (weighted) total latency. Again, a negative result shows that very simple symmetric load balancing games on identical machines do not admit better than 2-mixed-efficient taxes. On the positive side, 2-mixed-efficient refundable taxes with respect to the (weighted) total latency can be computed in polynomial time using convex quadratic programming. The case of nonrefundable taxes is more difficult to handle. However, there are load balancing games where ρ-mixed-efficient nonrefundable taxes can be computed for values of ρ which are smaller than the price of anarchy of the original game (without taxes). The interested reader can refer to [16] and [17] for the related results, proofs, and open problems.

9.4 Multicast Cost Sharing Games

In this section, we consider multicast cost sharing games. A *cost sharing method* for multicast games is a function $\mathscr{M}$ which, given a set of receivers R and a strategy profile σ, distributes among the receivers the total cost $cost(\Pi(\sigma))$ in such a way that $\sum_{i \in R} \mathscr{M}(R,\sigma,i) = cost(\Pi(\sigma))$, where $\mathscr{M}(R,\sigma,i) =_{def} \omega_i(\sigma)$ is the cost charged to receiver i.

Several cost sharing methods have been proposed so far, namely,

- $\mathscr{M}_1$ (*egalitarian*) equally shares the global cost among all the receivers, i.e.,

$$\mathscr{M}_1(R,\sigma,i) = \frac{cost(\Pi(\sigma))}{n}.$$

- $\mathscr{M}_2$ (*path-proportional*) shares the cost of each link $e \in \Pi(\sigma)$ among all the receivers j using it proportionally to the overall cost of their chosen path, i.e.,

$$\mathscr{M}_2(R,\sigma,i) = \sum_{e \in \Pi(\sigma_i)} c_e \frac{cost(\Pi(\sigma_i))}{\sum_{j: e \in \Pi(\sigma_j)} cost(\Pi(\sigma_j))}.$$

- $\mathscr{M}_3$ (*egalitarian-path-proportional*) shares the overall cost among all the receivers proportionally to the cost of their chosen path, i.e.,

$$\mathscr{M}_3(R,\sigma,i) = cost(\Pi(\sigma)) \frac{cost(\Pi(\sigma_i))}{\sum_{j \in R} cost(\Pi(\sigma_j))}.$$

- $\mathscr{M}_4$ (*Shapley*) equally shares the cost of each link $e \in \Pi(\sigma)$ among all the receivers using it, i.e.,

$$\mathscr{M}_4(R,\sigma,i) = \sum_{e \in \Pi(\sigma_i)} \frac{c_e}{n_e(\sigma)}.$$

Since cost sharing games naturally arise in socioeconomic scenarios, cost sharing methods are usually required to meet some constraining properties. The ones most standard and used are:

- *Weak budget balance*. A receiver is never charged a cost share greater than the cost of her chosen path, that is, $\mathscr{M}(R,\sigma,i) \leq cost(\Pi(\sigma_i))$.
- *Strong budget balance*. The cost of each link is only shared among the receivers using it.
- *Stability*. The game induced by $\mathscr{M}$ possesses pure Nash equilibria.
- *Fairness*. The cost share charged to any two receivers adopting two equivalent paths in a given strategy profile is the same, where two paths are equivalent if they have the same cost and are shared in the same way by paths having the same cost. More formally, let $a_k(e,\sigma)$ be the set of paths of cost k using edge e in σ; two paths p and p' are equivalent in σ if $cost(p) = cost(p')$ and $\sum_{e \in p \cap a_k(e,\sigma)} c_e = \sum_{e \in p' \cap a_k(e,\sigma)} c_e$ for any $k \geq 0$.
- *Separability*. The cost shares of each edge are computed independently and are completely determined by the set of receivers using it.

We call feasible any method meeting weak budget balance and fairness. Clearly, the strong budget balance property implies the weak one. In [20] it is remarked that the only method meeting all such properties is the Shapley value. Moreover, it is not difficult to see that the path proportional one meets strong budget balance and fairness, while the egalitarian-path-proportional one meets weak budget balance and fairness. Hence, all the above methods besides the egalitarian one are feasible.

Directed graphs. In [8] it is shown that all methods except for the path-proportional one meet stability. The price of anarchy for the game yielded by the egalitarian method is unbounded, while for all the other ones it is equal to n with respect to two different social cost functions: the overall transmission cost (function γ_{sum}), which coincides with the sum of all the shared costs, and the maximum shared cost paid by the receivers (function γ_{max}). As for the price of stability, in [2] it is proved that it is $\Theta(\log n)$ for the Shapley value with respect to γ_{sum}.

For any feasible method, the number of best response moves needed to reach a Nash equilibrium starting from any strategy profile can be arbitrarily large, even when $n = 2$. Motivated by these results, it becomes interesting to evaluate the price of anarchy after a limited number of best responses. For the egalitarian and path-proportional methods the price of anarchy is unbounded for any sequence of best responses and one-round walks. For the Shapley value method, the price of anarchy after a one-round walk is $\Theta(n^2)$. Such a value is outperformed by the egalitarian-path-proportional method, which achieves a price of anarchy of $\mathcal{O}(n)$ after a one-round walk. This is an asymptotically optimal result since it can be shown that any feasible method cannot achieve a price of anarchy better than n after a one-round walk. All the above results have been determined for both γ_{sum} and γ_{max} in [8].

For k-round walks, [26] shows that, starting from any strategy profile, the Shapley method achieves a price of anarchy of at most $\mathcal{O}(n\sqrt{n})$ after two rounds and gives a general lower bound of $\Omega(n\sqrt[k]{n})$ for any number $k \geq 2$ of rounds. Similarly, when starting from the empty strategy profile, exactly matching upper and lower bounds equal to n are determined for any number of rounds. When starting from an arbitrary strategy profile, both the egalitarian and the path-proportional methods can yield an unbounded price of anarchy, while the egalitarian path-proportional achieves a price of anarchy of $\Theta(n)$. Finally, when starting from the empty strategy profile, all these three methods achieve a price of anarchy of $\Theta(n)$.

Undirected graphs. We briefly describe results related to undirected graphs by outlining the differences with the directed case. When not explicitly claimed, all the presented results are taken from [8]. Deciding whether the path-proportional method meets stability is still an open problem. The price of anarchy of the egalitarian path-proportional method falls between $\frac{2}{3}n$ and n. Asymptotic lower bounds equal to 1.915, 2, and 1.714 are known, respectively, for the price of stability of the path-proportional, the egalitarian path-proportional, and the Shapley value methods with respect to γ_{sum} [7]. The price of anarchy of the egalitarian path-proportional method after a one-round walk at least $\frac{2}{3}n$. No general results on the performance of feasible methods are known. The main differences with the case of directed graphs hold for the best-response walks. For the price of anarchy of the Shapley value, only the trivial lower bound of $\Omega(n)$ is known for k-round walks when starting from any arbitrary strategy profile. When starting from the empty strategy profile, in [18] an upper bound of $\mathcal{O}(\log^3 n)$ and a lower bound of $\Omega(\log n)$ on the price of anarchy achievable after any number of rounds is proved. For one-round walks, an improved $\Omega(\log^2 n)$ lower bound is derived. For the egalitarian method we have a price of anarchy of $\Theta(\log n)$, for the path-proportional one we have a lower bound of $\Omega(\log n)$ and an upper bound of $\mathcal{O}(n)$, and for the egalitarian path-proportional a lower bound of $\Omega(n/\log n)$ and an upper bound of $\mathcal{O}(n)$.

Finally, finding a Nash equilibrium minimizing the potential function is NP-hard [19], as is finding the sequence of the best responses leading to the lowest possible social cost even after a one-round walk for both γ_{sum} and γ_{max}.

9.5 Communication Games in All-Optical Networks

All-optical networks have been largely investigated in recent years due to the promise of data transmission rates several orders of magnitude higher than those of current networks. The key to high speeds in all-optical networks is to maintain the signal in optical form, thereby avoiding the prohibitive overhead of conversion to and from electrical form at the intermediate nodes. The high bandwidth of the optical fiber is utilized through *wavelength-division multiplexing*: two signals connecting different source-destination pairs may share a link, provided they are transmitted on carriers having different wavelengths (or colors) of light. Since the optical spectrum is a scarce resource, a given communication pattern in optical networks is often designed so as to minimize the total number of colors used, a measure which is trivially lower-bounded by the maximum load, that is, the maximum number of connecting paths sharing the same physical edge.

In [12], the following problem is investigated. An all-optical network provider must determine suitable payment functions for non-cooperative agents wishing to communicate so as to induce Nash equilibria using a low number of wavelengths. More formally, an all-optical network is modeled as an undirected graph $G=(V,E)$ where nodes in V represent sites and undirected edges in E represent bidirectional optical fiber links between the sites. A point-to-point communication requires us to establish a uniquely colored path between the two nodes whose color is different from the colors of all the other paths sharing some of its edges. Each player in the game asks for the implementation of a certain point-to-point communication and is charged a cost by the network provider obtained by applying a certain payment function. Depending on the variables defining the payment function, it is possible to distinguish among three different levels of information needed by an agent in order to compute her cost charge.

- *Minimal*. Each agent i knows all the wavelengths available along any path in S_i.
- *Intermediate*. Each agent i knows the wavelengths available along any edge in the network.
- *Complete*. Each agent i knows the whole strategy profile.

Under the complete level, suitable payment functions can be computed such that in any Nash equilibrium the assignment of paths and colors to the agents is the same as with the one computed by a centralized algorithm aiming to minimize the optical spectrum, and the computational complexity is the same as that of the algorithm. For the remaining two levels, the most reasonable payment functions either do not admit pure Nash equilibria or induce games having the worst possible price of anarchy, that is, always possessing a pure Nash equilibrium assigning a different color to each agent. However, by suitably restricting the network topology, a price of anarchy of 8 has been obtained for chains and 16 for rings under the minimal level, and further reduced respectively to 3 and 6 under the intermediate level, up to an additive factor converging to 0 as the load increases. Finally, again under the minimal level, a price of anarchy logarithmic in the number of agents has been determined for trees.

9.6 Beyond Nash Equilibria: An Alternative Solution Concept for Non-cooperative Games

As witnessed by the related notions of price of anarchy and stability, pure Nash equilibria usually generates suboptimal solutions for non-cooperative games. One of the several reasons for this situation is the fact that players always perform selfish moves only motivated by a transient improvement on their payoffs, without considering what will be their final payoffs when the game eventually reaches a pure Nash equilibrium. This observation naturally yields the question of whether agents taking decisions only based on what will be their short-term consequences, without considering what these decisions will cause tomorrow, might be considered as rational.

In [11], Bilò and Flammini propose a new definition of selfish agents by giving them a more farsighted view of their actions. An agent knows she is part of a multiplayer game and also knows that the game will not stop right after she has performed an improving step. Assume that in state σ player i has an improving step and that if she performs such a move a sequence of improvements begins leading the game toward a pure Nash equilibrium σ'. In this scenario, player i is mostly interested in comparing the payoff she is experiencing in σ with the one she can get at σ'. The idea is that, if $\omega_i(\sigma')$ is worse than $\omega_i(\sigma)$, player i is damaged by the consequences of her improving step; hence, she had better not performed it. When more than just one equilibrium can be reached from a particular state, by following a classical worst-case analysis, it can be assumed that the agent will compare the payoff in the current state with that at the equilibrium yielding the worst payoff for her. Such a viewpoint is clearly based upon the definition of a ground set of equilibria the agents will compare a generic state with. According to these comparisons, a possible nonempty set of new equilibria may arise and the process may be iterated recursively until a fixed point is reached and the final set of desired equilibria is created. Such a set is called the set of Second Order equilibria. Using different definitions of equilibrium for defining the ground set, it is possible to achieve different sets of Second Order equilibria. [11] focuses on the definition and the evaluation of Second Order Nash equilibria, that is, with a ground set given by the set of pure Nash equilibria.

Consider the following generalization of the Nash dynamics graph $D(\mathscr{G})$. Given a set of (equilibria) states $E \subseteq S$, let $D(\mathscr{G},E) = (N,A)$ be a directed graph in which $N = S$ and there exists an edge between σ and σ' if and only if there exists an improving step from σ to σ' and $\sigma \notin E$. Edges are still labeled with the index of the player performing the related improving step. Clearly, $D(\mathscr{G},\emptyset)$ coincides with the Nash dynamics graph of $\mathscr{G}$. Define $\rho_E(\sigma)_i^k$ as the set of all the states of $\mathscr{G}$ that can be reached starting from σ by following a path of length at most k whose first arc is labeled with index i in the graph $D(\mathscr{G},E)$. The set $\rho_E(\sigma)^k = \bigcup_{i=1}^n \rho_E(\sigma)_i^k$ will denote the set of all states that can be reached from σ by following a path of length at most k in $D(\mathscr{G},E)$. When $E = \emptyset$, we will simply remove the subscript E from the notation. Let us also define $P(\sigma)$ as the set of players having an improving step in σ. Assume that the payoffs of each player are costs to be minimized. [11] introduces the following definition.

Definition 9.1. Let $\mathscr{G}$ be a game with the FIP property. The set $N^k(\mathscr{G}) = \{\sigma \in S : \forall i \in P(s)$ and $\forall \sigma' \in \rho(\sigma)_i^1, \exists \sigma'' \in N^k(\mathscr{G})$ such that $\sigma'' \in \rho_{N^k(\mathscr{G})}(\sigma')^k$ and $\omega_i(\sigma) < \omega_i(\sigma'')\}$ is the set of all the Second Order k-Nash equilibria of $\mathscr{G}$, for any integer $k \geq 0$.

Intuitively, this rather involved definition says that a state σ is a Second Order k-Nash equilibrium, for some integer $k \geq 0$, if all the players that have an improving step in σ would experience a payoff worse than the one they get in σ in one of the Second Order k-Nash equilibria resulting from an evolutive process of at most k improving steps taking place after their first defection. Such a definition is clearly recursive. However, it can be shown that it is well posed, in the sense that it admits a unique set of solutions or fixed points. First of all, it is easy to see that $N^0(\mathscr{G})$ coincides with the set of pure Nash equilibria for $\mathscr{G}$ and that each pure Nash equilibrium is a Second Order k-Nash equilibrium, for any integer $k \geq 1$. Then, by proving that there exists a value k^* for which all the sets $N^k(\mathscr{G})$ become the same for any $k \geq k^*$, the following final definition can be given.

Definition 9.2. Let $\mathscr{G}$ be a game with the FIP property. Each state $\sigma \in N^{k^*}(\mathscr{G}) =_{def} N(\mathscr{G})$ is a Second Order Nash equilibrium.

Clearly, since $N^0(\mathscr{G}) \subseteq N^k(\mathscr{G})$ for any $k \geq 0$, we have that the price of stability of Second Order Nash equilibria is not worse than that of pure Nash equilibria, while the price of anarchy of pure Nash equilibria is not worse than that of Second Order Nash equilibria. The interested reader is referred to [11] for applications of Second Order Nash equilibria to some traditional games, as well as extensions and variations of this notion of equilibrium.

9.7 Coping with Incomplete Information

In highly decentralized and pervasive networks, the assumption that each agent knows the strategy adopted by any other agent may be too optimistic or even infeasible. In such situations, the set of agents of which each agent knows the chosen strategy is modeled by means of a social knowledge graph, that is, a directed graph $K = (P, E)$ whose set of nodes coincides with the set of players in the game, and there is a directed edge from player i to player j if and only if i knows the strategy adopted by j. Since i always knows her chosen strategy, it is assumed that any social knowledge graph contains all self-loops. For each game $\mathscr{G}$ and social knowledge graph K, a graphical game $(\mathscr{G}, K)$ is obtained by extending its definition in such a way that the payoff of each player can be influenced only by the strategies adopted by the adjacent ones. In the following, we discuss graphical linear congestion games and graphical Shapley cost sharing games which have been studied in [9] and [10] respectively.

Graphical Linear Congestion Games. In a graphical congestion game $(\mathscr{G}, K)$, the payoff of player i is defined as $\omega_i(\sigma) = \sum_{e \in \sigma_i} d_e(n_e^i(\sigma, K))$, where $n_e^i(\sigma, K) = |\{ j \in P : e \in \sigma_j \text{ and } (i, j) \in E(K)\}|$ is the number of players using e in σ that are neighbors of i in K, including i herself.

The classical potential function defined by Rosenthal no longer applies to the graphical case; hence, as a first approach to the problem, a complete characterization of the cases possessing pure Nash equilibria and the FIP property needs to be achieved. The topology of K plays a fundamental role in this issue. In particular, if K is undirected, $(\mathscr{G}, K)$ is an exact potential game and thus isomorphic to a classical congestion game; if K is directed, an equilibrium is not guaranteed to exist in general, but $(\mathscr{G}, K)$ possesses the FIP property and an equilibrium can be found in polynomial time if K is acyclic, even if finding the best equilibrium remains an intractable problem.

Limited knowledge of the players yields two different possible definitions for the latencies experienced by a player in each state: The *presumed latency* is the one the player believes she suffers due to the fact that she is only aware of the existence of her neighbors, and is defined as $pres_i(\sigma) = \omega_i(\sigma)$. The *perceived latency* is the one she actually experiences due to all the players using the resource, and is defined as $perc_i(\sigma) = \sum_{e \in \sigma_i} d_e(n_e(\sigma))$. According to these definitions, four different social functions can be asked to be minimized, namely, the sum and the maximum of the presumed latencies and the sum and the maximum of perceived ones. Given a bound Δ on the maximum degree of K, for all the cases in which pure Nash equilibria are guaranteed to exist, tight lower and upper bounds on the price of stability and asymptotically tight bounds on the price of anarchy of pure strategies have been achieved for such social functions. These results have been also extended to load balancing games and are summarized in Table 9.1.

Graphical Shapley Cost Sharing Games. In a graphical Shapley cost sharing game $(\mathscr{G}, K)$, the payoff of player i is defined as $\omega_i(\sigma) = \sum_{e \in \sigma_i} \frac{c_e}{n_e^i(\sigma, K)}$. Again, because

Table 9.1 Summary of the results on the price of anarchy and stability with respect to the presumed and perceived social cost functions

(a) Presumed Latency

Presumed Latency	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}
Undirected graph	2, $\Theta(\Delta+1)$	$\Theta(\Delta+1)$, $\Delta+1$
Acyclic DAG	$\Theta(\Delta+1)$, $\Delta+1$	$\Theta(\Delta+1)$, $\Delta+1$

(b) Perceived Latency in Congestion Games

Perceived Latency	Congestion Games	
	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}
Undirected graph	n, $n \div n\sqrt{\Delta+1}$	$\Theta(n(\Delta+1))$
Acyclic DAG	$\Theta(n(\Delta+1))$	$\Theta(n(\Delta+1))$

(c) Perceived Latency in Load Balancing Games

Perceived Latency	Load Balancing Games	
	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}
Undirected graph	n, $\Theta(n)$	$\Theta(n)$
Acyclic DAG	$\Theta(n)$	$\Theta(n)$

of the fact that some receivers can be hidden to other ones, graphical Shapley cost sharing games can no longer be isomorphic to congestion games when considering the presence of social knowledge graphs.

If K is a directed acyclic graph (DAG), the same technique used for graphical linear congestion games shows that $(\mathscr{G}, K)$ possesses the FIP property and that an equilibrium can be computed in polynomial time. If K is either undirected or directed cyclic, existence of pure Nash equilibria is no longer guaranteed even for the multicast case. However, when K is undirected, the restriction to the load balancing case can be shown to be isomorphic to general potential games, and hence to have the FIP property. This does not hold when K is a directed graph containing cycles.

The results on the price of anarchy and stability which can be achieved on the multicast case are quite surprising. The price of stability achievable by any DAG is at least $\frac{1}{2}\log n$, while the price of anarchy for complete DAGs can be shown to be at most $\log^2 n$. This result can be achieved by proving that the set of Nash equilibria induced by any complete DAG K^* on any instance I coincides with the set of solutions obtained after a first round of best responses in which the receivers enter sequentially the game I starting from the empty configuration according to their topological ordering in K^*. The upper bound on the price of stability follows by exploiting a result presented in [18]. Putting everything together, we have that the complete DAG, if used as a universal knowledge graph, is able to contain the price of anarchy of the graphical Shapley multicast cost sharing game under a polylogarithmic bound.

When a particular instance of the Shapley multicast cost sharing game is fixed in advance, we have that the price of stability achieved by any DAG must be at least $\frac{4n}{n+3}$. On the other side, it is possible to prove that for any instance I there always exists a DAG $K(I)$ achieving a price of anarchy of at most $\frac{4n}{n+3}$ if $n = 2, 3$ and of at most $\frac{4(n-1)}{n+1}$ if $n \geq 4$, hence obtaining an upper bound on the price of anarchy almost matching the lower bound on the price of stability achievable by any DAG. Unfortunately, it is not known how to construct efficiently the graph $K(I)$. However, given any r-approximation of the optimal multicasting tree, it is possible to compute in polynomial time (by using a simple depth-first search) a DAG achieving a price of anarchy of at most $\frac{4n}{n+3}r$ if $n = 2,3$ and of at most $\frac{4(n-1)}{n+1}r$ if $n \geq 4$, lowering the price of anarchy of this game to a constant value.

These achievements are twofold: from one side, they shed some light on how the lack of knowledge among players can have impact on the total cost of the self-emerging networks created by the interactions of selfish users; from the other side, they show that the idea of hiding some players to others is a powerful instrument that a central authority managing the game can use in order to obtain solutions whose cost may be not too far from the optimal one without interfering directly on the choices made by the players. This situation can be seen as another evidence of the famous Braess' paradox [13]. According to this paradox, there are cases in which adding fast links in a network results in a decrease of performance or, symmetrically, hiding some fast links from the players can improve the network performance. This naturally extends to graphical Shapley multicast cost sharing games where hiding some of the players to other ones can yield better solutions. In this sense, the less players know, the most they are "cooperative."

9.8 Open Problems and Future Research

Algorithmic Game Theory is a relatively young research area. As a consequence, it provides a wide landscape for setting new ideas, theories, and applications. Also, lots of questions have been raised by the research conducted so far. In this section we propose a list of the most important and interesting open problems related to the topics covered in the chapter.

Congestion games. The price of anarchy and stability of linear congestion games have been exactly estimated both in the general case and in the load balancing case. Much of the current trend in the study of congestion games is thus moving towards the characterization of more general cases allowing more complicated delay functions as well as different demands; see, for instance, [1, 30]. Most of these results are concerned with the price of anarchy when considering, as a social function, the sum of the latencies experienced by all the players, while the study of the price of stability and the analysis of the social function given by the maximum delay experienced by the players still deserves further investigation. Also, problems related to the existence of pure Nash equilibria, the speed of convergence towards Nash equilibria, and the complexity of computing Nash equilibria in various special cases deserve further investigation.

Multicast cost sharing games. This is a widely open research topic since one can define several reasonable socioeconomic properties a cost sharing method should satisfy. A first effort in this direction can be found in [7, 20]. The latter paper, in particular, provides a complete characterization of the price of anarchy which can be achieved by cost sharing methods satisfying certain desiderata. Such properties are rather strong and a similar study based on weaker constraints like the ones defined in [7] could be an interesting research direction. Moreover, fewer results are known regarding the price of stability. In this particular setting, the major open problem regards the characterization of the price of stability of the Shapley value method in undirected networks, which, in spite of a constant lower bound, has been upper-bounded by $\mathcal{O}(\log n)$. A sublogarithmic upper bound is only known for the special case of broadcasting games [28].

Communication games in all-optical networks. While the complete information level has been fully understood, the main question left open under the lower levels is the determination of functions that yield Nash equilibria on every topology with performance better than the worst possible one assigning a different color to each agent. Moreover, under incomplete information, it would be interesting to improve and extend the results on specific topologies. In this latter case, moreover, nice connections between payment functions and online algorithms, allowing us to cope with the arbitrary order of the moves of the agents, have been shown. It would be challenging to understand the conditions and eventual systematic methods that could yield converging payment functions preserving online algorithm performance under incomplete information.

Second-order Nash equilibria. A lot of open questions have been introduced by considering the extended notion of rationality of selfish players. The first one is certainly that of giving further validation to Second Order equilibria by using them in conjunction with other known equilibria notions and presenting good applications. To this end, the definition of Second Order Sink equilibria seems to be a promising research direction. Moreover, there is the important issue of understanding the power of different ordering strategies in influencing the performance of these equilibria. An interesting question can be also that of trying to understand if the use of Second Order equilibria can lead sequences of improving steps towards better states. In [11] only impatient prudent agents have been considered. A final open issue is certainly that of analyzing the other three possible definitions for rational agents as well as the Second Order equilibria yielded by rush agents.

Coping with incomplete information. Possible applications of social knowledge graphs include the design of protocols and P2P systems which limit the visibility of the other peers, or simply, at a more foundational level, the possibility of using them as an intermediate methodological tool for defining cost shares and payoffs so as to induce good overall performance without direct interference in user decisions.

For Graphical linear congestion games, a crucial observation is that better bounds with respect to the ones reported in Table 9.1 can be obtained for specific social graphs. In fact, for the undirected complete graph, constant bounds are derived directly from the classical linear congestion game. Many questions are left open. Besides tightening the various constant multiplicative gaps, it would be worth closing the gap between the upper and the lower bound on the price of stability for the maximum perceived latency social function. The investigation of the consequences of a minimum degree very close to n is another interesting issue. Moreover, what about nonlinear latency functions? While the convergence for directed acyclic graphs works for all the latency functions, in the undirected case, convergence strictly relies on linearity. It would be also worth investigating the expected price of stability and anarchy when the knowledge graph obeys some social behavior, as in Kleinberg's Small Word model [33, 34]. In general, are there universal bounded degree social graphs always guaranteeing good performance?

For Graphical multicast cost sharing games, besides closing the gaps between the upper and lower bounds on the price of anarchy and stability (such as the gap for the price of anarchy in the universal case), it would be interesting to extend the study to other graphical cost sharing games, like cost sharing congestion games. Finally, what about the effect of social graphs on the speed of convergence, that is, on the number of selfish moves needed to reach equilibria or on the performance achieved after a limited number of steps?

Acknowledgements This work was partially supported by the European Union under the IST FET Integrated Project AEOLUS and EU COST action 293 – Graphs and Algorithms in Communication Networks (GRAAL) – and by a "Caratheodory" research grant from the University of Patras.

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. In: Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS '06), LNCS 3884, pp. 218-229, 2006
2. Anshelevich, E., Dasgupta, A., Kleinberg, J. M., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS '04), IEEE Computer Society, pp. 295-304, 2004
3. Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T.: Near-optimal network design with selfish agents. In: Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC '03). ACM Press, pp. 511-520, 2003
4. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC '05), ACM Press, pp. 57-66, 2005
5. Azar Y., Epstein, A.: The hardness of network design for unsplittable flow with selfish users. In: Proceedings of the 3rd Workshop on Approximation and Online Algorithms (WAOA '05), LNCS 3879, pp. 41-54, 2005
6. Azar, Y., Jain, K., Mirrokni, V. S.: (Almost) optimal coordination mechanisms for unrelated machine scheduling. In: Proceedings of the 19th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA '08), pp. 323-332, 2008

7. Bilò, V.: The price of Nash equilibria in multicast transmission games. In: Proceedings of the 18th International Symposium on Algorithms and Computation (ISAAC '07), LNCS 4835, pp. 390-401, 2007
8. Bilò, V., Fanelli, A., Flammini, M., Melideo, G., Moscardelli, L.: Designing fast converging cost sharing methods for multicast transmissions. Theory of Computing Systems, online first, 2009
9. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Graphical congestion games with linear latencies. In: Proceedings of the 20th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '08), pp. 194-196, 2008
10. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: When ignorance helps: graphical multicast cost sharing games. In: Proceedings of the 32nd International Symposium on Mathematical Foundations of Computer Science (MFCS '08), LNCS 5162, pp. 108-119, 2008
11. Bilò, V., Flammini, M.: Extending the notion of rationality of selfish agents: second order Nash equilibria. In: Proceedings of the 32nd International Symposium on Mathematical Foundations of Computer Science (MFCS '07), LNCS 4708, pp. 621-632, 2007
12. Bilò, V., Flammini, M., Moscardelli, L.: On Nash equilibria in non-cooperative all-optical networks. In: Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS '05), LNCS 3404, pp. 448-459, 2005
13. Braess, D.: Über ein Paradoxon der Verkehrsplanung. Unternehmensforschung, 12:258-268, 1968
14. Caragiannis, I.: Efficient coordination mechanisms for unrelated machine scheduling. In: Proceedings of the 20th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA '09), pp. 815–824, 2009
15. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., and Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06), LNCS 4051, Part I, pp. 311-322, 2006
16. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Taxes for linear atomic congestion games. In: Proceedings of the 14th Annual European Symposium on Algorithms (ESA '06), LNCS 4168, pp. 184-195, 2006
17. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Improving the efficiency of load balancing games through taxes. In: Proceedings of the 4th International Workshop on Internet and Network Economics (WINE '08), LNCS 5385, pp. 374-385, 2008
18. Charikar, M., Mattieu, C., Karloff, H., Naor, J., Saks, M.: Best response dynamics in multicast cost sharing. In: Proceedings of the 20th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '08), pp. 70-76, 2008
19. Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., Orda, A.: Non-cooperative multicast and facility location games. In: Proceedings of the 7th ACM Conference on Electronic Commerce (EC '06), ACM Press, pp. 72-81, 2006
20. Chen, H., Roughgarden, T., Valiant, G.: Designing networks with good equilibria. In: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '08), ACM Press, pp. 854-863, 2008
21. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC '05), ACM Press, pp. 67-73, 2005
22. Christodoulou, G., Koutsoupias, E.: On the price of anarchy and stability of correlated equilibria of linear congestion games. In Proceedongs of the 13th Annual European Symposium on Algorithms (ESA '05), LNCS 3669, pp. 59-70, 2005
23. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. In: Proceedings of the 31st Annual International Colloquium on Automata, Languages, and Programming (ICALP '04), LNCS 3142, pp. 345-357, 2004
24. Christodoulou, G., Mirrokni, V., Sidiropoulos, A.: Convergence and approximation in potential games. In: Proceedings of the 23rd Symposium on Theoretical Aspects of Computer Science (STACS '06), LNCS 3884, Springer, pp. 349-260, 2006

25. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC ’03), ACM Press, pp. 521-530, 2003
26. Fanelli, A., Flammini, M., Moscardelli, L.: On the convergence of multicast games in directed networks. In: Proceedings of the 19th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA ’07), ACM Press, pp. 330-338, 2007
27. Fanelli, A., Flammini, M., Moscardelli, L.: The speed of convergence in congestion games under best-response dynamics. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP ’08), LNCS 5125. Part I, pp. 796-807, 2008
28. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the price of stability for designing undirected networks with fair cost allocations. In: Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP ’06), LNCS 4051, Part I, pp. 608-618, 2006
29. Fleischer, L., Jain, K., Mahdian, M.: 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 ’04), IEEE Computer Society, pp. 277-285, 2004
30. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish unsplittable flows. Theoretical Computer Science, 348(2-3): 226-239, 2005
31. Immorlica, N., Li, L., Mirrokni, V. S., Schulz, A.: Coordination mechanisms for selfish scheduling. In: Proceedings of the 1st Workshop on Internet and Network Economics (WINE ’05), LNCS 3828, pp. 55-69, 2005
32. Karakostas, G., Kolliopoulos, S.: Edge pricing of multicommodity networks for heterogeneous selfish users. In: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS ’04), IEEE Computer Society, pp. 268-276, 2004
33. Kleinberg, J.: The small-world phenomenon: an algorithm perspective. In: Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC ’00), ACM Press, pp. 163-170, 2000
34. Kleinberg, J.: Small-world phenomena and the dynamics of information. In: Proceedings of the 14th Advances in Neural Information Processing Systems (NIPS ’01), pp. 431-438, 2001
35. Koutsoupias, E., Papadimitriou, C. H.: Worst-case Equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS ’99), LNCS 1653, pp. 404-413, 1999
36. Mirrokni, V. S., Vetta, A.: Convergence issues in competitive games. In: Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX ’04), LNCS 3122, pp. 183-194, 2004
37. Monderer, D., Shapley, L. S.: Potential games. Games and Economic Behavior. 14:124-143, 1996
38. Nash, J.: Equilibrium points in *n*-person games. In: Proceedings of the National Academy of Sciences, volume 36, pp. 48-49, 1950
39. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. V.: Algorithmic game theory. Cambridge University Press, 2007
40. Pigou, A. C.: The economics of welfare. Macmillan, 1920
41. Rosenthal, R. W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory, 2:65-67, 1973
42. Roughgarden, T.: Stackelberg scheduling strategies. SIAM Journal on Computing, 33(2): 332-350, 2004
43. Roughgarden, T.: On the severity of Braess’s Paradox: designing networks for selfish users is hard. Journal of Computer and System Sciences, 72(5): 922-953, 2006
44. Roughgarden, T., Tardos, E.: How bad is selfish routing? Journal of the ACM, 49(2): 236-259, 2002
45. Shapley, L. S.: The value of *n*-person games. Contributions to the theory of games, pp. 31-40, Princeton University Press, 1953
46. Suri, S., Tóth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. Algorithmica, 47(1): 79-96, 2007