Game Theoretic Study of Networks

George Christodoulou*

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
gchristo@di.uoa.gr

Abstract. In this thesis we study the efficiency of systems, in which, users share resources. We assume that the users are selfish and we use principles of Game Theory for our study. In the first part of the thesis, we theoretically study the influence of the systems efficiency due to selfishness. The basic tool that we use in our study is the Price of Anarchy [32, 44]. In the second part of the thesis, we study algorithms (mechanisms), in order to remedy the situation due to selfishness. A Mechanism is an algorithm that aims at leading the players into actions desirable for the system and achieve this goal by appropriately modifying the parameters of the game.

1 Introduction

During the last decade, the rapid growth of the Internet, lead the Computer Science community into a new effort for understanding some of its most critical features. Internet consists of a variety of autonomous computational entities that react into the system, having different and often conflicting interests. Non cooperative Game Theory, studies situations (games), where selfish agents with conflicting interests, take decisions ignoring the social welfare. This observation motivated the creation of a new scientific area that lies in the border between Computer Science and Economics and particularly Game Theory. The object of this new area is the modeling of a distributed system as a (non) cooperative game, where the reactions taking place are strategic decisions of the selfish autonomous entities (agents, players) that constitute the network.

Having such a model at hand, two main questions that motivated our research are the following:

– How much is the system's performance affected due to the users' selfish behavior?
– Can we modify the parameters of the system, so as to implicitly enforce the selfish agents to act in a way that benefits the social welfare?

This paper consists of two parts that cover a wide spectrum of the Game Theoretic perspective of network analysis. The first part (Section 2) concerns a theoretical study of the influence of the systems' performance due to selfishness. The main tool that we use in our study is the Price of Anarchy [32, 44]. In

* Dissertation Advisor: Elias Koutsoupias, Professor

the second part (Section 3), we try to eliminate the undesirable situations that appear due to selfishness, by use of mechanisms. A mechanism is an algorithm that aims at enforcing the agents to react according to the systems' designer desire and achieve this by appropriately modifying the parameters of the game.

2 The Price of Anarchy of Congestion Games

In this section we study the influence of the systems performance due to selfishness. We consider a wide and interesting subclass of games introduced by Rosenthal, naming Congestion Games[45].

Related Work The study of the price of anarchy was initiated in [32], for (weighted) congestion games of m parallel links. The price of anarchy for the maximum social cost is proved to be $\Omega(\frac{\log m}{\log \log m})$, while in [38, 33, 18] they proved $\Theta(\frac{\log m}{\log \log m})$. In [18], they extended the result to m parallel links with different speeds and showed that the price of anarchy is $\Theta(\frac{\log m}{\log \log \log m})$. In [17], more general latency functions are studied, especially in relation to queuing theory. For the same model of parallel links, [25] and [37] consider the price of anarchy for other social costs.

In [50], the special case of congestion games in which each strategy is a singleton set is considered. They give bounds for the case of the average social cost. For the same class of congestion games and the maximum social cost, [26] showed that the price of anarchy is $\Theta(\log N/ \log \log N)$ (a similar, perhaps unpublished, result was obtained by the group of [50]). The case of singleton strategies is also considered in [27] and [37]. In [24], they consider the mixed price of anarchy of symmetric network weighted congestion games, when the network is layered.

The non-atomic case of congestion games was considered in [46, 47] where they showed that for linear latencies the average price of anarchy is 4/3. They also extended this result to polynomial latencies.

2.1 The Pure Price of Anarchy of Congestion Games

In [11], we study the price of anarchy of *pure equilibria* in general congestion games with linear latency functions.

We consider both the maximum and the average (sum) player cost as social cost. We also study both symmetric and asymmetric games. Our results (both lower and upper bounds) are summarized in Table 1. For the case of asymmetric games, the values hold also for network congestion games. We don't know if this is true for the symmetric case as well.

We extend these results to the case of latency functions that are polynomials of degree p with nonnegative coefficients. The results (both lower and upper bounds) appear Table 2.

2.2 The Correlated Price of Anarchy and of Stability

In [10], we study linear general congestion games with linear cost (latency) functions. We focus mainly on the sum social cost which is the sum of the cost of

	SUM	MAX
Symmetric	$\frac{5N-2}{2N+1}$	$\frac{5N+1}{2N+2} \dots \frac{5}{2}$
Asymmetric	$\frac{5}{2}$	$\Theta(\sqrt{N})$

Table 1. Price of anarchy of pure equilibria for linear latencies. N is the number of the players.

	SUM	MAX
Symmetric	$p^{\Theta(p)}$	$p^{\Theta(p)}$
Asymmetric	$p^{\Theta(p)}$	$\Omega(N^{p/(p+1)}) \dots O(N)$

Table 2. Price of anarchy of pure equilibria for polynomial latencies of degree p. N is the number of the players.

all players and we consider all types of equilibria: dominant strategies, pure and mixed Nash equilibria, and correlated equilibria.

These equilibria are related by inclusion and this hierarchy allows us to look for the strongest possible results. In particular, when we obtain a lower bound on the price of stability or the price of anarchy for dominant strategies, this lower bound holds for all types of equilibria. (It is important to emphasize that this holds for the price of stability because a strong dominant strategy implies unique Nash and correlated equilibrium). And on the other end, when we obtain an upper bound for correlated equilibria, this holds for all types of equilibria. Interestingly—but not entirely unexpectantly—such general results are easier to prove in some cases (when we are not distracted by the additional structure of the specific subproblems).

Price of stability: For linear congestion games we give an upper bound of 1.6 for Nash and correlated equilibria. Although bounding directly the price of stability seems hard—after all, we need to bound the best not the worst equilibrium—we resort to a clever trick by making use of the potential of congestion games. More specifically, instead of bounding the cost of the best equilibrium, we bound the cost of the pure Nash equilibrium which has minimum potential. Since every local optimum of the potential corresponds to a pure Nash equilibrium (by the handy theorem of Rosenthal [45, 40]), such a Nash equilibrium is guaranteed to exist.

We give a non-trivial lower bound of $1 + \frac{\sqrt{3}}{3} \approx 1.577$ for dominant strategies. This is a surprisingly strong result: It states in the strongest possible way that for some games, selfishness deteriorates the efficiency of some systems by approximately 58%. It is also perhaps the most technical part of this work. Naturally, both the upper and lower bounds hold for all types of equilibria (dominant strategies, pure and mixed Nash, and correlated equilibria). We also observe that for the max social cost (i.e., the maximum cost among the players) the price of stability is $\Theta(\sqrt{N})$.

Price of anarchy: For linear congestion games, we extend some of the results of our STOC'05 paper [11] on the price of anarchy. There we showed bounds on Nash equilibria which we extend here to correlated equilibria. At the same time we strengthen the bounds. More specifically, we show that the correlated price of anarchy of the sum social cost is 2.5 for the asymmetric case and $\frac{5N-2}{2N+1}$ for the symmetric case, where N is the number of players. Since in [11], we had matching lower bounds for pure Nash equilibria, these are also tight bounds for pure and mixed Nash and correlated equilibria.

We also extend the results of [6] about the price of anarchy of Nash equilibria for weighted linear congestion games when the social cost is the total latency: The price of anarchy of correlated equilibria is $\frac{3+\sqrt{5}}{2} \approx 2.618$. Although we prove a more general result, our proof is substantially simpler.

2.3 Convergence in Congestion Games

The main tool for analyzing the performance of systems where selfish players interact without central coordination, is the notion of the *price of anarchy* in a game; this is the worst case ratio between an optimal social solution and a Nash equilibrium. Intuitively, a high price of anarchy indicates that the system under consideration requires central regulation to achieve good performance. On the other hand, a low price of anarchy does not necessarily imply high performance of the system. One main reason for this phenomenon is that in many games, the repeated selfish behavior of players may not lead to a Nash equilibrium. Moreover, even if the selfish behavior of players converges to a Nash equilibrium, the *rate* of convergence might be very slow. Thus, from a practical and computational viewpoint, it is important to evaluate the rate of convergence to approximate solutions.

By modeling the repeated selfish behavior of the players as a sequence of atomic improvements, the resulting convergence question is related to the running time of local search algorithms. In fact, the theory of PLS-completeness [49] and the existence of exponentially long walks in many local optimization problems such as Max-2SAT and Max-Cut, indicate that in many of these settings, we cannot hope for a polynomial-time convergence to a Nash equilibrium. Therefore, for such games, it is not sufficient to just study the value of the social function at Nash equilibria. To deal with this issue, we need to bound the social value of a strategy profile after *polynomially many* best-response improvements by players.

Our Contribution Our work [14] deviates from bounding the distance to a Nash equilibrium [49, 21], and focuses in studying the rate of convergence to an approximate solution [39, 28]. We consider two types of walks of best responses: random walks and deterministic fair walks. On random walks, we choose a random player at each step. On deterministic fair walks, the time complexity of a game is measured in terms of the number of *rounds*, where a round consists of a sequence of movements, with each player appearing at least once in each round.

First, we give tight bounds for the approximation factor of the solution after one round of best responses of players in selfish routing games. In particular, we prove that starting from an arbitrary state, the approximation factor after one round of best responses of players is at most $O(n)$ of the optimum and this is tight up to a constant factor. We extend the lower bound for the case of multiple rounds, where we show that for any constant number of rounds t, the approximation guarantee cannot be better than $n^{\epsilon(t)}$, for some $\epsilon(t) > 0$. On the other hand, we show that starting from an empty state, the state resulting after one round of best responses is a constant-factor approximation.

3 Mechanisms

As was pointed out in the previous section, equilibrium points in general lead to solutions that are not optimal in a social sense and sometimes the distance from the optimum is high *(Price of Anarchy is high)*.

In this section we will concentrate on techniques that attempt to resolve this problem. The common in these approaches is that they all aim at altering the players objectives, in a way that bad equilibria of the game are eliminated.

3.1 Coordination Mechanisms

Here, we propose an algorithmic framework in order to reduce the price of anarchy, naming *Coordination Mechanisms*.

Related Work Mechanisms to improve coordination of selfish agents is not a new idea and we only mention here work that directly relates to our approach. A central topic in game theory [43] is the notion of mechanism design in which the players are paid (or penalized) to "coordinate". The differences between mechanism design and the coordination mechanism model are numerous. The most straightforward comparison can be exhibited in the selfish routing problem: both aim at improving coordination, but mechanism design can be seen as a way to introduce *tolls* (see for example [15, 16]), while coordination mechanism is a way to introduce *traffic lights*.

Cole, Dodis and Roughgarden [15] considered an heterogeneous variation of the problem. They associate to each user a a scalar $\alpha(a)$ that indicates the user's sensitivity to the taxation. They show that for any single-commodity network with heterogeneous users there always exist price vectors that result in a minimum latency Wardrop equilibrium. Their result is existential and is proved with a fixed-point theorem (see also [22], for a different proof that holds also for the multiple source, single sink setting). Using some additional assumptions on the heterogeneity function α and on the network latency functions they show how to compute the lower possible prices that lead to a minimum latency equilibrium.

Fleischer et. al [23], and Karakostas and Kolliopoulos [30] independently showed that optimal taxes exist for the multi-commodity case as well. Their proof relies on standard tools mathematical programming. They also prove that the optimal taxes are efficiently computable if the optimum flow is given. For the homogeneous case ($\alpha(a) = 1$), Cole, Dodis and Roughgarden [16] show that in every network with linear latencies, taxes cannot improve the cost of a flow at Nash. In this case the benefit of taxes, in the best case equals the benefit of edge removals. If non linear latencies are considered, then taxes may be more powerful than edge removal. Karakostas and Kolliopoulos [31], show that for linear latencies, the social cost at equilibrium with taxation is at most twice the social optimum cost without taxation and at most 5/4 times the social optimum with taxation.

Contribution We study coordination mechanisms for congestion games [12]. We show an interesting relation between the potential and the social cost of a set of strategies; based on these we give a coordination mechanism with price of anarchy n for the single-commodity congestion games. We also show that the bound n is tight. We conjecture that the same bound holds for the general congestion games; but we were able to show only that the coordination mechanism that we employed for the single-commodity games fails in the general case.

3.2 Mechanism Design

The study of mechanisms in game-theoretic settings is an important area at the intersection of Computer Science and Game Theory. A particular type of mechanisms, for which auctions is a typical example, is the mechanism design problem. Mechanisms are a special class of algorithms and the study of their computational properties was initiated by Nisan and Ronen in their seminal paper [41]. The focus of their paper was on the task allocation problem on unrelated machines. They showed that no mechanism can have approximation ratio better than 2. They conjectured that this lower bound is not tight. In this paper, we confirm this and improve the lower bound of the approximation ratio to $1 + \sqrt{2}$.

The problem is one of the most fundamental scheduling problems [29, 36]. There are n machines and m tasks and each task may have different execution times on the machines. Let t_{ij} be execution time of task j on machine i. The objective is to schedule the tasks on the machines to minimize the makespan. In the mechanism design setting, each machine i knows its own times (the t_{ij}'s), but the algorithm does not know them. The algorithm first asks the machines to declare their times t_{ij} and then proceeds to allocate the tasks according to a policy known to machines in advance. The machines are selfish players who are lazy and don't want to execute the tasks, so they may lie. To deal with this problem, the mechanism pays the machines according to their declarations. Thus the mechanism design problem consists of two algorithms: an allocation algorithm and a payment algorithm. They take as input the declaration of times by the machines and produce an allocation and a set of payments, one for each machine.

The objective of each machine is to minimize the load of tasks allocated to it, minus its payment. On the other hand, the objective of the mechanism is to minimize the makespan of the allocation. Notice that the mechanism does not care how much he pays to the machines. The payments are given to machines as an incentive to tell the truth. A mechanism is called *truthful* when telling the truth is a dominant strategy for each player, independently of the declarations of the other players.

There are two major classes of problems in algorithmic mechanism design. For every problem of the first class, there exists an optimal truthful mechanism but the problem is NP-hard (i.e., the problem of computing the optimal allocation is NP-hard). For this kind of problems, we are interested on truthful *polynomial-time approximation* algorithms. Two typical problems in this class are the problem of combinatorial auction and the problem of scheduling related machines. The second class contains problems that need not be NP-hard, but for which no optimal mechanism is truthful. The quintessential problem in this class is the scheduling problem. For this kind of problems, we can ask either about the optimal approximation ratio of *all* algorithms, or the optimal approximation ratio of *polynomial-time* algorithms. In this paper, we deal with the approximation ratio of all algorithms, not necessarily polynomial-time ones. In other words, *the lower bound of* $1+\sqrt{2}$ *is based on the restrictions imposed only by truthfulness*, not by the computational hardness of the problem.

Related Work

The scheduling problem on unrelated machines is one of the most fundamental scheduling problems [29, 36]. Here we study its mechanism design version and we improve the results of Nisan and Ronen [41, 42], who introduced the problem and initiated the algorithmic theory of Mechanism Design. They gave a truthful n-approximate (polynomial-time) algorithm; they also showed that no mechanism (polynomial-time or not) can achieve approximation ratio better than 2. They conjectured that there is no deterministic mechanism with approximation ratio less than n. On the other hand, they gave a randomized truthful mechanism for two players, that achieves an approximation ratio of 7/4.

Archer and Tardos [5] considered the variant of the problem for related machines. In this case, for each machine there is a single value (instead of a vector), its speed. They provided a characterization of all truthful algorithms in this class, in terms of a monotonicity condition. Using this characterization, they gave a variant of the optimal algorithm which is truthful (albeit exponential-time). They also gave a polynomial-time randomized 3-approximation mechanism, which was later improved to a 2-approximation, in [3]. This mechanism is truthful in expectation. Andelman, Azar, and Sorani [1] gave a 5-approximation deterministic truthful mechanism, in the same framework, which was later improved by Kovacs[34] to a 3-approximation.

Much more work has been done in the context of combinatorial auctions (see for example [4, 8, 9, 19, 7, 20] and the references within). In this setting, Saks and Yu [48] proved that for convex domains the Monotonicity Property characterizes the class of truthful mechanisms, generalizing results of [28, 35].

Our Contribution

In [13] we give the first lower bound greater than 2 for the mechanism design version of the task scheduling problem. In particular we prove that there is no deterministic truthful mechanism that can achieve an approximation ratio of less than $1+\sqrt{2}$ for an instance of the problem with 3 or more machines.

In order to prove the lower bound, we consider an instance of the problem with 3 machines and 5 tasks. We combine three properties that every truthful mechanism must satisfy and that are implications of the Monotonicity Property.

The first property is due to Nisan and Ronen [42]. They have used it to obtain their lower bounds in their original paper. It is a specific and direct way to take advantage of the Monotonicity Property. It states that if a machine gets a set of tasks when he declares t_i, it will get exactly the same set of tasks if we lower the execution time of the tasks allocated to the machine and increase the execution time of the remaining tasks. The second is a useful 2-dimensional property of truthful mechanisms. It states that if we fix all the players except for one, the way that the allocation procedure partitions the space of that player, depending on his bidding values, cannot be arbitrary, but has a particular shape. The third property holds only for three machines. It states that under particular circumstances, if we know that the allocation for some specific instances that differ only in the bidding values of a single player, remain the same, then if we alter all the players bidding values together, then the allocation also remains the same.

An appropriate combination of the aforementioned properties leads to a construction of an instance in which either a player gets all the tasks or the tasks are assigned to machines with high prosecution times. In both cases, a truthful mechanism cannot escape from the approximation ratio of $1 + \sqrt{2}$.

References

1. N. Andelman, Y. Azar, and M. Sorani. Truthful approximation mechanisms for scheduling selfish related machines. In *Proc. of the 22nd*, LNCS, pages 69–82, 2005.
2. Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *45th Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.
3. Aaron Archer. *Mechanisms for Discrete Optimization with Rational Agents.* PhD thesis, Cornell University, January 2004.
4. Aaron Archer, Christos H. Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 205–214, 2003.
5. Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *42nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 482–491, 2001.
6. Baruch Awerbuch, Yossi Azar, and Amir Epstein. Large the price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.
7. Moshe Babaioff, Ron Lavi, and Elan Pavlov. Mechanism design for single-value domains. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference (AAAI)*, pages 241–247, 2005.
8. Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *Submitted to SODA07*, 2006.
9. P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanims design. In *Proc. of 37th*, 2005.
10. George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Algorithms - ESA 2005, 13th Annual European Symposium*, pages 59–70, 2005.
11. George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.
12. George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. In *Automata, Languages and Programming: 31st International Colloquium (ICALP)*, pages 345–357, 2004.

13. George Christodoulou, Elias Koutsoupias, and Angelina Vidali. A lower bound for scheduling mechanisms. In *(Accepted) ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.
14. George Christodoulou, Vahab S. Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. In *23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 349–360, 2006.
15. Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 521–530, 2003.
16. Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.
17. Artur Czumaj, Piotr Krysta, and Berthold Vöcking. Selfish traffic allocation for server farms. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing (STOC).*, pages 287–296, 2002.
18. Artur Czumaj and Berthold Vöcking. Tight bounds for worst-case equilibria. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 413–420, 2002.
19. Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 610–618, 2005.
20. Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 644–652, 2006.
21. A. Fabrikant, C. Papadimitriou, and K. Talwar. On the complexity of pure equilibria. In *STOC*, 2004.
22. Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. In *Automata, Languages and Programming: 31st International Colloquium (ICALP)*, pages 544–554, 2004.
23. Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *45th Symposium on Foundations of Computer Science (FOCS)*, pages 277–285, 2004.
24. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. In *ICALP*, 2004.
25. Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. Computing nash equilibria for scheduling on restricted parallel links. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 613–622, 2004.
26. Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. The price of anarchy for polynomial social cost. In *Mathematical Foundations of Computer Science 2004, 29th International Symposium (MFCS)*, pages 574–585, 2004.
27. Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. Nash equilibria in discrete routing games with convex latency functions. In *Automata, Languages and Programming: 31st International Colloquium (ICALP)*, pages 645–657, 2004.
28. M. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS*, page to appear, 2005.
29. D.S. Hochbaum. *Approximation algorithms for NP-hard problems.* PWS Publishing Co. Boston, MA, USA, 1996.

30. George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *45th Symposium on Foundations of Computer Science (FOCS)*, pages 268–276, 2004.
31. George Karakostas and Stavros G. Kolliopoulos. The efficiency of optimal taxes. In *Combinatorial and Algorithmic Aspects of Networking, First Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN)*, pages 3–12, 2004.
32. E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
33. Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. Approximate equilibria and ball fusion. *Theory Comput. Syst.*, 36(6):683–693, 2003.
34. Annamária Kovács. Fast monotone 3-approximation algorithm for scheduling related machines. In *Algorithms - ESA 2005, 13th Annual European Symposium*, pages 616–627, 2005.
35. R. Lavi, A. Mu'Alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. of*, 2003.
36. J.K. Lenstra, D.B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46(1):259–271, 1990.
37. Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. A new model for selfish routing. In *21st Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 547–558, 2004.
38. Marios Mavronicolas and Paul G. Spirakis. The price of selfish routing. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 510–519, 2001.
39. V.S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *RANDOM-APPROX*, pages 183–194, 2004.
40. D. Monderer and L. Shapley. Potential games. *Games and Economics Behavior*, 14:124–143, 1996.
41. Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing (STOC)*, pages 129–140, 1999.
42. Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.
43. Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory.* MIT Press, 1994.
44. Christos H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
45. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
46. T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
47. T. Roughgarden and E. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.
48. Michael E. Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings 6th ACM Conference on Electronic Commerce (EC)*, pages 286–293, 2005.
49. A. Schaffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.
50. Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Selfish load balancing and atomic congestion games. In *Proceedings of the Sixteenth Annual ACM Symposium on Parallel Algorithms (SPAA)*, pages 188–195, 2004.