Mediated Equilibria in Load-Balancing Games*

Joshua R. Davis, David Liben-Nowell, Alexa Sharp, and Tom Wexler

Carleton College; Northfield, MN Oberlin College; Oberlin, OH
joshuadavis@q.com, dlibenno@carleton.edu, {alexa.sharp,tom.wexler}@oberlin.edu

Abstract. *Mediators* are third parties to whom the players in a game can delegate the task of choosing a strategy; a mediator forms a *mediated equilibrium* if delegating is a best response for all players. Mediated equilibria have more power to achieve outcomes with high social welfare than Nash or correlated equilibria, but less power than a fully centralized authority. Here we begin the study of the power of mediation by using the mediation analogue of the price of stability—the ratio of the social cost of the best mediated equilibrium BME to that of the socially optimal outcome OPT. We focus on load-balancing games with social cost measured by weighted average latency. Even in this restricted class of games, BME can range from as good as OPT to no better than the best correlated equilibrium. In unweighted games BME achieves OPT; the weighted case is more subtle. Our main results are (1) that the worst-case ratio BME/OPT is at least $(1 + \sqrt{2})/2 \approx 1.2071$ (and at most $1 + \phi \approx 2.618$ [3]) for linear-latency weighted load-balancing games, and that the lower bound is tight when there are two players; and (2) tight bounds on the worst-case BME/OPT for general-latency weighted load-balancing games. We also give similarly detailed results for other natural social-cost functions.

1 Introduction

The recent interest in algorithmic game theory by computer scientists is in large part motivated by the recognition that the implicit assumptions of traditional algorithm design are ill-suited to many real-world settings. Algorithms are typically designed to generate solutions that can be implemented by some centralized authority. But often no such centralized authority exists; solutions arise through the interactions of self-interested, independent agents. Thus researchers have begun to use game theory to model these competitive, decentralized situations.

One classic example is the paper of Koutsoupias and Papadimitriou [16], who consider the effect of decentralizing a standard load-balancing problem. In the resulting game, each job is controlled by a distinct player who selects a machine to serve her job so as to minimize delay. The authors compare the social cost (expected maximum delay) of the Nash equilibria of this game to that of a centrally designed optimal solution. The maximum of these ratios is the *price of anarchy* of the game, quantifying the worst-case cost of decentralized behavior.

* All authors were supported in part by NSF grant CCF-0728779. This work was also supported by grants from Carleton College, Oberlin College, and Denison University. Thanks to Bobby Kleinberg for suggesting this line of inquiry to us.

One can imagine a continuum indexing the amount of power that a centralized authority has in implementing solutions to a given problem, from utter impotence (leading to a potentially inefficient Nash equilibrium) to dictatorial control (leading to the socially optimal outcome). For example, consider a weak authority who can *propose* a solution simultaneously to all players, but who has no power to enforce it. The players would agree to such a proposal only if it were a Nash equilibrium—but the authority could propose the *best* Nash equilibrium. The ratio of the cost of the best Nash equilibrium to the global optimum is the *price of stability*, which may be much better than the price of anarchy.

A *correlator* is a more powerful authority, in that it is not required to broadcast the entire proposed solution; it signals each player individually with a suggested action, chosen from some known joint probability distribution. The resulting stable outcomes are called *correlated equilibria* [2]. Any Nash equilibrium is a correlated equilibrium, but often a correlator can induce much better outcomes.

A *mediator* [1, 19, 21–24, 26] is an authority who offers to act on behalf of the players; any player may *delegate* to the mediator the responsibility of choosing a strategy. In a *mediated equilibrium*, all players prefer to delegate than to play on their own behalf. The strategies that the mediator selects for the delegating players may be correlated; moreover, the distribution from which the mediator draws these strategies may depend on which players have opted to delegate. A mediator can enforce an equilibrium by threatening to have the delegating players "punish" any player departing from mediation. Any correlated equilibrium can be represented as a mediated equilibrium, but the converse is not true; mediators are more powerful than correlators.

The present work: mediated load-balancing games. In this paper, we begin to quantify the powers and limitations of mediators. We consider the mediation analogue of the price of stability: how much less efficient than the globally optimal outcome OPT is the best mediated equilibrium BME? (While one could ask questions analogous to the price of anarchy instead, the spirit here is that of a well-intentioned central authority who would aim for the best, not the worst, outcome within its power.) We initiate this study in the context of *load-balancing games*. Each player i controls a job that must be assigned to a machine. Each machine j has a nonnegative, nondecreasing latency function $f_j(x)$, and each player incurs a *cost* of $f_j(\ell_j)$ for choosing machine j, where ℓ_j is the total load of jobs on machine j. We split load-balancing games into classes along two dimensions:

- *unweighted* vs. *weighted*: in weighted games, job i has weight w_i and experiences cost $f_j(\sum_{i' \text{ uses } j} w_{i'})$ on machine j; in unweighted games all $w_i = 1$.
- *linear* vs. *general* latencies: in linear games, $f_j(x) = a_j \cdot x$ for $a_j \geq 0$; for general latencies f_j can be an arbitrary nonnegative, nondecreasing function.

The social cost is measured by the weighted average latency experienced by the jobs; see Section 6 for results using other social cost functions.

Load-balancing games are appealing for this work for two reasons. First, they include cases in which mediators can achieve OPT and cases in which they cannot even better the best Nash equilibrium BNE. Second, the prices of anarchy and

	unweighted jobs	**weighted jobs**
linear latencies	BME = OPT [19] BCE $\leq 4/3 \cdot$ BME [tight] (Lemma 1 [4])	BME $\leq 2.618 \cdot$ OPT, BCE $\leq 2.618 \cdot$ BME [3] $n = 2$: BME $\leq 1.2071 \cdot$ OPT (Thm. 2) BCE $\leq 4/3 \cdot$ BME (Thm. 2) [both tight for $n = 2$]
general latencies	BME = OPT [19] BCE $\leq n \cdot$ BME [tight] (Lemma 1)	BME $\leq \Delta \cdot$ OPT [tight] (Thm. 3) BCE $\leq \Delta \cdot$ BME [tight] (Thm. 3)

Fig. 1. Summary of our results for weighted-average-latency social cost. Here OPT is the socially optimal outcome, BME (BCE) the best mediated (correlated) equilibrium, n the number of jobs, and Δ the ratio of total job weight to smallest job weight.

stability, and corresponding measures of correlated equilibria, are well understood for these games and many of their variants [3–6, 15–17, 25]. Most relevant for what follows are an upper bound of $1 + \phi \approx 2.618$ on the price of anarchy in weighted linear games [3] and a tight upper bound of 4/3 on the price of stability in unweighted linear games [4].

We extend this line of work to mediated equilibria with the following results. (Figure 1 summarizes those for the weighted-average-latency social cost.)

- In the unweighted case, the BME is optimal, regardless of the latency functions' form. This result follows from a recent theorem of Monderer and Tennenholtz [19], which in fact holds for any symmetric game. See Section 3.
- In weighted linear-latency games with two players, we give tight bounds on the best solution a mediator can guarantee: a factor of $(1 + \sqrt{2})/2 \approx 1.2071$ worse than OPT but 4/3 better than the best correlated equilibrium BCE. Thus mediators lie strictly between dictators and correlators. See Section 4.
- In weighted nonlinear-latency games, mediated equilibria provide no worst-case improvement over correlated or even Nash equilibria. See Section 5.
- We also analyze mediation under two other social cost functions that have been considered in the literature: (i) the maximum latency of the jobs; and (ii) the average latency, unweighted by the jobs' weights. See Section 6.

Related work. Koutsoupias and Papadimitriou initiated the study of the price of anarchy in load-balancing games, considering weighted players, linear latencies, and the maximum (rather than average) social cost function [16]. A substantial body of follow-up work has improved and generalized their initial results [6, 7, 9, 18]. See [12] and [27] for surveys. A second line of work takes social cost to be the sum of players' costs. Lücking et al. [11, 17] measure the price of anarchy of mixed equilibria in linear and convex routing games in this setting. Awerbuch et al. [3] consider both the unweighted and weighted cases on general networks. Suri et al. [15, 25] examine the effects of asymmetry in these games. Caragiannis et al. [4] give improved bounds on the price of anarchy and stability. Christodoulou and Koutsoupias [5, 6] bound the best- and worst-case correlated equilibria in addition to improving existing price of anarchy and stability results.

Other aspects of correlated equilibria have been explored recently, including their existence [13] and computation [13, 14, 20]. Mediated equilibria have developed in the game theory literature over time; see Tennenholtz [26] for a summary. Mediated equilibria have been studied for position auctions [1], for network routing games [23, 24], and in the context of social choice and voting [21, 22]. Strong mediated equilibria have also been considered [19, 24].

2 Notation and Background

An *n-player, m-machine load-balancing game* is defined by a nondecreasing *latency function* $f_j : [0, \infty) \to [0, \infty]$ for each machine $j \in \{1, \ldots, m\}$; and a *weight* $w_i > 0$ for each player $i \in \{1, \ldots, n\}$. We consider games in which every job has access to every machine: a pure strategy profile $\mathbf{s} = \langle s_1, \ldots, s_n \rangle$ can be any element of $\mathcal{S} := \{1, \ldots, m\}^n$. The *load* ℓ_j on a machine j under $\mathbf{s}$ is $\sum_{i:s_i=j} w_i$, and the *latency* of machine j is $f_j(\ell_j)$. The *cost* $c_i(\mathbf{s})$ to player i under $\mathbf{s}$ is $f_{s_i}(\ell_{s_i})$. Pure Nash equilibria exist in all load-balancing games [8, 10]. A load-balancing game is *linear* if each f_j is of the form $f_j(x) = a_j \cdot x$ for some $a_j \geq 0$ and *unweighted* if each $w_i = 1$. Machine j is *dominated* by machine j' for player i if, no matter what machines the other $n - 1$ players use, player i's cost is lower using machine j' than using machine j.

A nonempty subset of the players is called a *coalition*. A *mediator* is a collection Ψ of probability distributions ψ_T for each coalition T, where the probability distribution ψ_T is over pure strategy profiles for the players in T. The *mediated game* M_Γ^Ψ is a new n-player game in which every player either participates in Γ directly by choosing a machine in $S := \{1, \ldots, m\}$ or participates by *delegating*. That is, the set of pure strategies in M_Γ^Ψ is $Z = S \cup \{s_{\text{med}}\}$. If the set of delegating players is T, then the mediator plays the correlated strategy ψ_T on behalf of the members of T. In other words, for a strategy profile $\mathbf{z} = \langle z_1, z_2, \ldots, z_n \rangle$ where $T := \{i : z_i = s_{\text{med}}\}$ and $\overline{T} := \{i : z_i \neq s_{\text{med}}\} = \{i : z_i \in S\} = \{1, \ldots, n\} - T$, the mediator chooses a strategy profile $\mathbf{s}_T$ according to the distribution ψ_T, and plays s_i on behalf of every player $i \in T$; meanwhile, each player i in $\overline{T}$ simply plays z_i. The expected cost to player i under the strategy profile $\mathbf{z}$ is then given by $c_i(\mathbf{z}) := \sum_{\mathbf{s}_T} c_i(\mathbf{s}_T, \mathbf{z}_{\overline{T}}) \cdot \psi_T(\mathbf{s}_T)$. (The mediators described here are called *minimal mediators* in [19], in contrast to a seemingly richer class that allow more communication from players to the mediator.)

A *mediated equilibrium* for Γ is a mediator Ψ such that the strategy profile $\langle s_{\text{med}}, s_{\text{med}}, \ldots, s_{\text{med}} \rangle$ is a pure Nash equilibrium in M_Γ^Ψ. Every probability distribution ψ' over the set of all pure strategy profiles for Γ naturally corresponds to a mediator Ψ, where the probability distribution ψ_T for a coalition T is the marginal distribution for T under ψ'—that is, $\psi_T(\mathbf{s}_T) = \sum_{\mathbf{s}' : \mathbf{s}'_T = \mathbf{s}_T} \psi'(\mathbf{s}')$. If ψ' is a correlated equilibrium then this Ψ is a mediated equilibrium.

The *social cost* of a strategy profile $\mathbf{s}$ is the total (or, equivalently, average) cost of the jobs under $\mathbf{s}$, weighted by their sizes—that is, $\sum_i w_i \cdot c_i(\mathbf{s})$. (We discuss other social cost functions in Section 6.) We denote by OPT the (cost of the) profile $\mathbf{s}$ that minimizes the social cost. We denote the worst Nash equilibrium—

the one that maximizes social cost—by WNE, and the best Nash (correlated, mediated) equilibrium by BNE (BCE, BME). Note that OPT $\leq$ BME $\leq$ BCE $\leq$ BNE $\leq$ WNE because every Nash equilibrium is a correlated equilibrium, etc. The *price of anarchy* is WNE/OPT, and the *price of stability* is BNE/OPT.

3 Unweighted Load-Balancing Games

Although the unweighted case turns out to have less interesting texture than the weighted version, we start with it because it is simpler and allows us to develop some intuition. We begin with an illustrative example:

Example 1. There are n unweighted jobs and two machines L and R with latency functions $f_L(x) = 1 + \varepsilon$ for any load, and $f_R(x) = 1$ for load $x > n - 1$ and $f_R(x) = 0$ otherwise.

For each player, R dominates L, so $\langle R, R, \ldots, R\rangle$ is the unique correlated and Nash equilibrium, with social cost n. Consider the following mediator Ψ. When all n players delegate, the mediator picks uniformly at random from the n strategy profiles in which exactly one player is assigned to L. When any other set of players delegates, those players are deterministically assigned to R. If all players delegate under Ψ, each player's expected cost is $(1 + \varepsilon)/n$; if any player deviates, then that player will incur cost at least 1. Thus Ψ is a mediated equilibrium. Its cost is only $1 + \varepsilon$, which is optimal, while BNE $=$ BCE $= n$.

In fact, this "randomize among social optima" technique generalizes to all unweighted load-balancing games—in any such game, BME $=$ OPT. This is a special case of a general theorem of Monderer and Tennenholtz [19] about mediated equilibria robust to deviations by coalitions. (See also [24].)

Example 1 shows that with nonlinear latency functions BCE may be much worse than OPT, even in the unweighted 2-machine case. But even linear unweighted load balancing has a gap between BCE and OPT, even in the 2-job, 2-machine case. The following example demonstrates the gap.

Example 2 (Caragiannis et al. [4]). There are two (unweighted) jobs and two machines L and R with latency functions $f_L(x) = x$ and $f_R(x) = (2 + \varepsilon) \cdot x$.

Here BCE $= 4$ and OPT $= 3 + \varepsilon$. (Machine L dominates R; no player can be induced to use R in any correlated equilibrium.) We can show that this example is tight with respect to the gap between BME and BCE, using a result on linear unweighted load-balancing games of Caragiannis et al. [4] and the "randomize among social optima" mediation technique. We can also show a tight bound for unweighted nonlinear latency load-balancing games (details omitted for space).

Lemma 1. *In n-player unweighted load-balancing games:*

- *for games with linear latency functions,* BCE $\leq 4/3 \cdot$ BME*. This bound is tight.*
- *for not-necessarily-linear latency functions,* BCE $\leq n \cdot$ BME*. This is tight.*

We now have a complete picture for unweighted load balancing: a tight bound on the gap between BME and BCE and the theorem that BME $=$ OPT.

4 Weighted Linear Load-Balancing Games

We now turn to weighted load-balancing games, where we find a richer landscape of results: among other things, cases in which BME falls strictly between BCE and OPT. We begin with the linear-latency case. (All proofs are omitted due to space.)

Theorem 2. *In any 2-machine, 2-job weighted game with linear latencies:*

1. BCE/BME $\leq 4/3$. *This bound is tight for an instance with weights* $\{1, 1\}$ *and with latency functions* $f_L(x) = x$ *and* $f_R(x) = (2 + \varepsilon) \cdot x$.
2. BME/OPT $\leq \frac{1+\sqrt{2}}{2}$. *This bound is tight for an instance with weights* $\{1, 1 + \sqrt{2}\}$ *and with latency functions* $f_L(x) = x$ *and* $f_R(x) = (1 + 2\sqrt{2}) \cdot x$.

The worst case for BCE/BME is actually unweighted—in fact, Example 2. This result fully resolves the 2-player, 2-machine case with linear latency functions. Included in this class of games are instances in which BCE > BME > OPT. One concrete example is with weights $\{1, 1 + \sqrt{2}\}$, $f_L(x) = x$, and $f_R(x) = \frac{3+3\sqrt{2}}{2} \cdot x$, when BCE/BME $= \frac{20+4\sqrt{2}}{23} \approx 1.1155$ and BME/OPT $= \frac{14\sqrt{2}-1}{17} \approx 1.1058$.

Adding additional machines to a 2-player instance does not substantively change the results (there is no point in either player using anything other than the two "best" machines), but the setting with $n \geq 3$ players requires further analysis, and, it appears, new techniques. Recent results on the price of anarchy in linear load-balancing games [3, 4, 6] imply an upper bound of $1 + \phi \approx 2.618$ on BME/OPT for any number of players n, where ϕ is the golden ratio. We believe that the worst-case ratio of BME/OPT does not decrease as n increases. (Consider an n-player instance in which $n - 2$ players have jobs of negligible weight and the remaining 2 players have jobs as in Theorem 2.) However, we do not have a proof that BME/OPT cannot worsen from $\frac{1+\sqrt{2}}{2} \approx 1.2071$ as n grows; nor do we have a 3-job example for which BME/OPT is worse than $\frac{1+\sqrt{2}}{2}$. The major open challenge emanating from our work is to close the gap between the upper bound (BME/OPT ≤ 2.618) and our bad example (BME/OPT $= 1.2071$) for general n.

5 Weighted Nonlinear Load-Balancing Games

We now consider weighted load-balancing games with latency functions that are not necessarily linear. We know from Lemma 1 that even in unweighted cases the power of Nash and correlated equilibria is limited. The weighted setting is even worse: the price of anarchy is unbounded, even if we restrict our attention to *pure* equilibria. Consider two identical machines, with latencies $f(x) = 0$ for $x \leq 5$ and $f(x) = 1$ for $x \geq 6$. There are four jobs, two of size 3 and two of size 2. A solution with cost zero exists (each machine has one size-2 and one size-3 job), but putting the two size-3 jobs on one machine and the two size-2 jobs on the other is a pure Nash equilibrium too. We can show that the price of stability is better in this setting, but in general BME is no better than BNE:

Theorem 3. *In any n-player weighted load-balancing game with job weights* $\{w_1, \ldots, w_n\}$ *(and not necessarily linear latency functions),* BNE $\leq \Delta \cdot$ OPT, *where* $\Delta := \sum_i w_i / \min_i w_i$ *is the ratio of total job weight to smallest job weight. Thus* BME $\leq \Delta \cdot$ OPT *and* BCE $\leq \Delta \cdot$ BME. *Both bounds are tight.*

6 Other Social-Cost Functions

Thus far we have discussed the social cost function $\mathsf{sc}_{wavg}(\mathbf{s}) := \sum_i w_i \cdot c_i(\mathbf{s})$ exclusively. Two other social cost functions have received attention in the literature: the maximum latency $\mathsf{sc}_{max}(\mathbf{s}) := \max_i c_i(\mathbf{s})$ and the unweighted average latency $\mathsf{sc}_{uavg}(\mathbf{s}) := \sum_i c_i(\mathbf{s})$. Under sc_{max}, in any load-balancing game BME = OPT = BNE: starting from OPT, run best-response dynamics (BRD) until it converges; no BRD step increases the maximum load, so the resulting Nash equilibrium is still OPT. Thus mediation is uninteresting under sc_{max}.

The behavior of sc_{uavg} turns out to be similar to that of sc_{wavg}. For nonlinear latencies, an analogue to Theorem 3 states that BCE $\leq n \cdot$ BME and BME $\leq n \cdot$ OPT (OPT is at least the maximum cost x experienced by a job in OPT; running BRD from OPT yields a Nash equilibrium where each job experiences cost at most x); the examples from Theorem 3 and Example 1 both remain tight. The 2-job linear case is also qualitatively similar; however, in contrast to the sc_{wavg} setting (where there is a bound of BME/OPT ≤ 2.618 for n-player games), even mediators cannot enforce outcomes that are close to OPT under sc_{uavg} as the number of players grows, even in linear-latency games. (Our construction also demonstrates that none of BCE, BNE, and WNE can provide constant approximations to OPT.)

Theorem 4. *Under* sc_{uavg}*, in linear-latency weighted load-balancing games:*

- *for 2 jobs and 2 machines,* BCE/BME $\leq \frac{4}{3}$ *(this is tight for Example 2) and* BME/OPT $\leq \frac{2+4\sqrt{2}}{7} \approx 1.0938$ *(this is tight for Theorem 2.2's example).*
- *for n jobs and 2 machines,* BME/OPT *is not bounded by any constant.*

7 Future Directions

In this paper we have begun to analyze the power of mediators in the spirit of price of stability, focusing on load-balancing games under the weighted average latency social cost function. We have a complete story for unweighted games and for weighted games with general latency functions. The biggest open question is the gap between BME and OPT in n-player weighted linear games. We know that for all such games BME/OPT ≤ 2.618 [3], and that there exist examples in which BME/OPT $= 1.2071$. What is the worst-case BME/OPT for $n \geq 3$ players? In particular, it may be helpful to understand better the connection between sc_{uavg} and sc_{wavg}: it was unexpected that the same instance is the worst case for both functions in the 2-player case (Theorem 2 and Theorem 4).

The broader direction for future research, of course, is to characterize the power of mediators in games beyond load balancing. For example, the upper bound of BME/OPT ≤ 2.618 in weighted linear load-balancing games comes from an upper bound on the price of anarchy in congestion games, a more general class of games. It is an interesting question as to how much better mediated equilibria are than correlated equilibria in, say, linear-latency weighted congestion games.

References

1. I. Ashlagi, D. Monderer, M. Tennenholtz. Mediators in position auctions. *EC*'07.
2. R. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.
3. B. Awerbuch, Y. Azar, A. Epstein. The price of routing unsplittable flow. *STOC*'05.
4. I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, L. Moscardelli. Tight bounds for selfish and greedy load balancing. *ICALP*'06.
5. G. Christodoulou, E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. *ESA*'05.
6. G. Christodoulou, E. Koutsoupias. The price of anarchy of finite congestion games. *STOC*'05.
7. A. Czumaj, B. Vöcking. Tight bounds for worst-case equilibria. *ACM Trans. on Algorithms*, 3(1):4, 2007.
8. E. Even-Dar, A. Kesselman, Y. Mansour. Convergence time to Nash equilibria. *ICALP*'03.
9. D. Fotakis, S. Kontogiannis, P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2):226–239, 2005.
10. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *ICALP*'02.
11. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, M. Rode. Nash equilibria in discrete routing games with convex latency functions. *J. Comput. Syst. Sci.*, 74(7):1199–1225, 2008.
12. M. Gairing, T. Lücking, B. Monien, K. Tiemann. Nash equilibria, the price of anarchy and the fully mixed Nash equilibrium conjecture. *ICALP*'05.
13. S. Hart, A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
14. S. Kakade, M. Kearns, J. Langford, L. Ortiz. Correlated equilibria in graphical games. *EC*'03.
15. A. Kothari, S. Suri, C. Tóth, Y. Zhou. Congestion games, load balancing, and price of anarchy. *Workshop on Combinatorial and Algorithmic Aspects of Networking*'04.
16. E. Koutsoupias, C. Papadimitriou. Worst-case equilibria. *STACS*'99.
17. T. Lücking, M. Mavronicolas, B. Monien, M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406(3):187–206, 2008.
18. M. Mavronicolas, P. Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007.
19. D. Monderer, M. Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009.
20. C. Papadimitriou, T. Roughgarden. Computing equilibria in multi-player games. *Journal of the ACM*, 55(3):1–29, 2008.
21. B. Peleg, A. Procaccia. Implementation by mediated equilibrium. Discussion paper 463, Center for the Study of Rationality, Hebrew University, 2007.
22. B. Peleg, A. Procaccia. Mediators and truthful voting. Working paper, 2008.
23. O. Rozenfeld, M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. *WINE*'06.
24. O. Rozenfeld, M. Tennenholtz. Routing mediators. *IJCAI*'07.
25. S. Suri, C. Tóth, Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.
26. M. Tennenholtz. Game-theoretic recommendations: Some progress in an uphill battle. *Proc. Intl. Conf. on Autonomous Agents and Multiagent Systems*'08.
27. B. Vöcking. Selfish load balancing. N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.