# Altruism in Atomic Congestion Games

MARTIN HOEFER, Max-Planck-Institut für Informatik
ALEXANDER SKOPALIK, University of Paderborn

This paper studies the effects of altruism, a phenomenon widely observed in practice, in the model of atomic congestion games. Altruistic behavior is modeled by a linear trade-off between selfish and social objectives. Our model can be embedded in the framework of congestion games with player-specific latency functions. Stable states are the pure Nash equilibria of these games, and we examine their existence and the convergence of sequential best-response dynamics. In general, pure Nash equilibria are often absent and existence is NP-hard to decide. Perhaps surprisingly, if all delay functions are affine, the games remain potential games even when agents are arbitrarily altruistic. The construction underlying this result can be extended to a class of general potential games and social cost functions, and we study a number of prominent examples. These results give important insights into the robustness of multi-agent systems with heterogeneous altruistic incentives. Furthermore, they yield a general technique to prove that stabilization is robust even with partly altruistic agents, which is of independent interest.

In addition to these results for uncoordinated dynamics, we consider a scenario with a central altruistic institution that can set incentives for the agents. We provide constructive and hardness results for finding the minimum number of altruists to stabilize an optimal congestion profile and more general mechanisms to incentivize agents to adopt favorable behavior. These results are closely related to Stackelberg routing and answer open questions raised recently in the literature.

Categories and Subject Descriptors: F.2.2 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]: Nonnumerical Algorithms and Problems — Computations on discrete structures; I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence — Multiagent systems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Congestion Games; Altruism; Potential Games; Nash equilibrium

**ACM Reference Format:**
Hoefer, M., Skopalik, A. 201x. Altruism in Atomic Congestion Games. ACM Trans. Econ. Comp. 1, 1, Article 1 (January 1), 21 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

## 1. INTRODUCTION

The study of algorithmic issues in systems with interacting rational agents has focused on models for a variety of important applications in the Internet, in traffic or social networks. A fundamental assumption in the analytical models, however, is that all

---

Some of the results in this paper have appeared in preliminary form as extended abstracts in the proceedings of ESA 2009 [Hoefer and Skopalik 2009a] and WINE 2009 [Hoefer and Skopalik 2009b].
This work was supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD), by DFG through UMIC Research Center at RWTH Aachen University, Cluster of Excellence MMCI at Saarland University, and grant Ho 3831/3-1, and by the German Israeli Foundation (GIF) under contract 877/05.
Author's addresses: M. Hoefer, Max-Planck-Institut für Informatik, Saarbrücken, Germany, mhoefer@mpi-inf.mpg.de; A. Skopalik, Department of Computer Science, University of Paderborn, Germany, skopalik@mail.uni-paderborn.de.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

agents are *selfish*. Their goals are restricted to optimizing their direct personal benefit, e.g. their personal delay in a routing game. The assumption of selfishness in the preferences of agents is found in the vast majority of present work on algorithmic aspects of economic problems in large networks. However, this assumption has been repeatedly questioned by economists and psychologists. In experiments it has been observed that participant behavior can be quite complex and contradictory to selfishness [Ledyard 1997; Levine 1998]. Various explanations have been given for this phenomenon, e.g., senses of fairness [Fehr and Schmidt 1999], reciprocity among agents [Gintis et al. 2005], or spite and altruism [Levine 1998; Eshel et al. 1998].

Prominent developments in the Internet like Wikipedia, open source software development, or Web 2.0 applications involve or explicitly rely on voluntary participation and contributions towards a joint project without direct personal benefit. These examples display forms of *altruism*, in which agents accept certain personal burdens (e.g., by investing time, attention, and money) to improve a common outcome. While malicious behavior has been considered recently for instance in nonatomic routing [Karakostas and Viglas 2007; Babaioff et al. 2009; Chen and Kempe 2008], virus inoculation [Moscibroda et al. 2009], or Bayesian congestion games [Gairing 2008], a deeper analysis of the effects of altruistic agents on competitive dynamics in algorithmic game theory is only starting to evolve.

We consider and analyze a model of altruism inspired by Ledyard [1997, p. 154], and studied for non-atomic routing games by Chen and Kempe [2008]. Each agent $i$ is assumed to be partly selfish and partly altruistic. Her incentive is to optimize a linear combination of personal cost and social cost, given by the sum of cost values of all agents. The strength of altruism of each agent $i$ is captured by her *altruism level* $\beta_i \in [0, 1]$, where $\beta_i = 0$ results in a purely selfish and $\beta_i = 1$ in a purely altruistic agent.

Chen and Kempe [2008] proved that in non-atomic routing games Nash equilibria are always guaranteed to exist, even for partially spiteful users, and analyzed the price of anarchy for parallel link networks. More recently, Chen et al. [2011] and Caragiannis et al. [2010] considered the price of anarchy in several classes of games with altruists. In our paper, we conduct the first study of the existence and complexity of equilibrium concepts with altruistic agents in atomic congestion games, a well-studied model for resource sharing. A standard congestion game is given by a set $N$ of myopic selfish users and a set $E$ of resources. Each resource $e$ has a non-decreasing delay function $d_e$. Every agent $i$ can pick a strategy $S_i$ from a set of possible strategies $\mathcal{S}_i \subseteq 2^E$, which means she allocates the set $S_i$ of resources (e.g. a path in a network). She then experiences a delay corresponding to the total delay on all resources in $S_i$, which in turn depends on the number of agents that allocate each resource. Each agent strives to pick a strategy minimizing her experienced delay. A stable state in such a game is a pure Nash equilibrium, in which each agent picks exactly one strategy, and no agent can decrease her delay by unilaterally changing her strategy. The study of congestion games received a lot of attention in recent years, mostly because of the intuitive formulation and their appealing analytical properties. In particular, they always possess a pure Nash equilibrium and every sequential better-response dynamics converges to one such equilibrium.

As one might expect, the presence of altruists can significantly alter the convergence and existence guarantees of pure Nash equilibria in congestion games. After a formal definition of *congestion games with altruists* in Section 2, we concentrate on pure equilibria and leave a study of mixed Nash equilibria for future work. Our results are as follows.

It is a simple exercise to observe that even in a *singleton game*, in which each strategy consists of a single resource, and for agents with *symmetric* strategy spaces, where

each agent has the same set of strategies, a pure Nash equilibrium can be absent. This is the case even for *pure altruists and egoists*, i.e. a population of agents which are either purely altruistic or purely selfish and their $\beta_i \in \{0, 1\}$. However, we show in Section 3 that such games admit a polynomial time algorithm to decide the existence problem. Furthermore, our algorithm can be adapted to compute the Nash equilibrium with best and worst social cost if it exists, for any agent population with a constant number of different altruism levels.

For slightly more general *asymmetric singleton games*, in which strategy spaces of agents differ, we show in Section 3 that deciding the existence of pure Nash equilibria becomes NP-hard. Nevertheless, for the important subclass of convex delay functions, i.e., linear and superlinear functions, previous results [Milchtaich 1996; Ackermann et al. 2009] imply that for any agent population a pure Nash equilibrium exists and can be obtained in polynomial time. In contrast, we show in Section 4.1 that convexity of delay functions is not sufficient for more general games. In particular, even for *symmetric network games*, in which strategies represent paths through a network, quadratic delay functions and pure altruists, pure Nash equilibria can be absent and deciding their existence is NP-hard.

Perhaps surprisingly, if all delay functions are *affine* with $d_e(x) = a_e x + b_e$, then there is a potential function. Thus, for every agent population pure Nash equilibria exist and better-response dynamics converges. We present this result in Section 4.2 in the form of a characterization for more general potential games. We outline a class of games and social cost functions that guarantee the existence of a potential function even upon introduction of altruists. As examples of such games we briefly study local interaction games and selfish scheduling with a Time-Sharing coordination mechanism. In Section 4.3 we extend to weighted congestion games and show that even for parallel links, identical delays $d_e(x) = x$ and pure altruists, pure Nash equilibria can again be absent and deciding their existence is NP-hard. If, however, we consider social cost as the *weighted* sum of cost values for all agents, then these games again are part of our class, and existence and convergence are guaranteed.

In addition to these results for uncoordinated dynamics, in Section 5 we consider a slightly more coordinated scenario, in which there is a central institution striving to obtain a good outcome. An obvious way to induce favorable behavior is to convince agents to act altruistically. In this context a natural question is how many altruists are required to stabilize a social optimum. This has been considered under the name "price of optimum" in [Kaporis and Spirakis 2009] for Stackelberg routing in nonatomic congestion games. As a Nash equilibrium in atomic games is not necessarily unique, we obtain two measures - an *optimal stability threshold*, which is the minimum number of altruists such that there is *any* optimal Nash equilibrium, and an *optimal anarchy threshold*, which asks for the minimum number of altruists such that *every* Nash equilibrium is optimal. For symmetric singleton games, we adapt our algorithm for computing Nash equilibria to determine both thresholds in polynomial time.

In our model the optimal anarchy threshold might not be well-defined even for singleton games. If all agents are altruists, there are suboptimal local optima in symmetric games with concave delays, or in asymmetric games with linear delays. Hence, even by making all agents altruists, the worst Nash equilibrium sometimes remains suboptimal. In contrast, we adapt the idea of the optimal stability threshold to a very general scenario, in which we can find a stable state with a given, not necessarily optimal, congestion profile. Each agent has a personalized *stability cost* for accepting a strategy under the given congestions. We provide an incentive compatible mechanism to determine an allocation of agents to strategies with minimum total stability cost. Unfortunately, such a general result is restricted to the case of singleton games. Even for symmetric network games on series-parallel graphs, we show that the problem of

determining the optimal stability threshold is NP-hard. Our reduction also yields inapproximability within any finite factor. This resolves an open problem raised by Kaporis and Spirakis [2008] on computing the "price of optimum" in atomic congestion games.

## 2. MODEL AND INITIAL RESULTS

A *congestion game with altruists* $G$ is given by a set $N$ of $n$ agents and a set $E$ of $m$ resources. Each agent $i$ has a set $\mathcal{S}_i \subseteq 2^E$ of strategies. In a *singleton* congestion game each agent has only singleton strategies $\mathcal{S}_i \subseteq E$. A vector of strategies $S = (S_1, \ldots, S_n)$ is called a *state*. Each agent $i$ has a personal weight $w_i$. We call a game *unweighted* if $w_1 = \ldots = w_n = 1$, and *weighted* otherwise. We consider weighted games only in Section 4.3. For brevity we use the term congestion game to refer to the unweighted version throughout and explicitly mention when weighted games are under consideration.

For a state we denote by $n_e = \sum_{i:e\in S_i} w_i$ the congestion, i.e. the total weight of agents using a resource $e$ in their strategy. Each resource $e$ has a *latency* or *delay* function $d_e(n_e)$, and the *delay for an agent* $i$ playing $S_i$ in state $S$ is $d_i(S) = \sum_{e\in S_i} d_e(n_e)$. The *social cost* of a state is the total delay of all agents $c(S) = \sum_{i\in N}\sum_{e\in S_i} d_e(n_e) = \sum_{e\in E} n_e d_e(n_e)$. Each agent $i$ has an *altruism level* of $\beta_i \in [0,1]$, and her *individual cost* is $c_i(S) = \beta_i c(S) + (1-\beta_i)d_i(S)$. We call an agent $i$ an *egoist* if $\beta_i = 0$ and a $\beta_i$*-altruist* otherwise. A *(pure) altruist* has $\beta_i = 1$, a *(pure) egoist* has $\beta_i = 0$. A game $G$ *with only pure altruists and egoists* is a game, in which $\beta_i \in \{0,1\}$ for all $i \in N$. A game $G$ is said to have $\beta$*-uniform* altruists if $\beta_i = \beta \in [0,1]$ for every agent $i \in N$. A (pure) *Nash equilibrium* is a state $S$, in which no agent $i$ can decrease her individual cost by unilaterally changing her strategy. We exclusively consider pure Nash equilibria in this paper.

With the exception of Section 4.3 we consider unweighted congestion games. If all agents are egoists, such ordinary (unweighted) congestion games have an exact potential function $\Phi(S) = \sum_{e\in E}\sum_{x=1}^{n_e} d_e(x)$ [Rosenthal 1973]. Thus, existence of Nash equilibria and convergence of iterative better-response dynamics are guaranteed. Obviously, if all agents are altruists, pure Nash equilibria correspond to local optima of the social cost function $c$ with respect to a local neighborhood consisting of single player strategy changes. Hence, existence and convergence are also guaranteed. This directly implies the same properties for $\beta$-uniform games, in which an exact potential function is $\Phi_\beta(S) = (1-\beta)\Phi(S) + \beta c(S)$.

In general, however, pure Nash equilibria might not exist.

PROPOSITION 2.1. *There are symmetric singleton congestion games with only pure altruists and egoists without a pure Nash equilibrium.*

*Example* 2.2. Consider a game with two resources $e$ and $f$, three egoists and one (pure) altruist. The delay functions are $d_e(x) = d_f(x)$ with $d_e(1) = 4$, $d_e(2) = 8$, $d_e(3) = 9$, and $d_e(4) = 11$. Then, in equilibrium each resource must be allocated by at least one egoist. In case there are two agents on each resource, the social cost is 32. In this case the altruist is motivated to change as the resulting cost is 31. In that case, however, one of the egoists on the resource with congestion 3 has an incentive to change. Thus, no pure Nash equilibrium will evolve.

Our interest is thus to characterize the games that have pure Nash equilibria. Towards this end we observe that an altruistic congestion game can be cast as a congestion game with player-specific latency functions [Milchtaich 1996]. In such a game the delay of resource $e$ to agent $i$ depends on the congestion and on the agent, i.e., $c_i(S) = \sum_{e\in S_i} d_e(n_e, i)$.

PROPOSITION 2.3. *A congestion game with altruists is a player-specific congestion game.*

PROOF. To embed our games within the framework of player-specific congestion games, we first consider a game with only pure altruists and egoists. An altruist moves from $S_i$ to $S_i'$ if the decrease in total delay $n_e d_e(n_e)$ on the resources $e \in S_i - S_i'$ she is leaving exceeds the increase on resources $e \in S_i' - S_i$ she is migrating to. Hence, altruists can be seen as myopic selfish agents with $c_i(S) = d_i'(S) = \sum_{e \in S_i} d_e'(n_e)$ with $d_e'(n_e) = n_e d_e(n_e) - (n_e - 1)d_e(n_e - 1)$, for $n_e > 0$. We set $d_e'(0) = 0$. This implies the result for games with pure altruists and egoists. Naturally, a $\beta_i$-altruist corresponds to a selfish agent with player-specific function $c_i(S) = (1 - \beta_i)d_i(S) + \beta_i d_i'(S)$. Thus, congestion games with altruists can be embedded into the class of player-specific congestion games. □

For some classes of player-specific congestion games it is known that pure Nash equilibria always exist. In particular, for games in which all individual delay functions are monotone and in which the strategy space of each agent is a matroid, existence and polynomial-time computation of a pure Nash equilibrium is guaranteed [Ackermann et al. 2009]. Non-existence in Example 2.2 is due to the fact that the individual delay function for the altruist is not monotone. Monotonicity holds, in particular, if delay functions are convex, which yields the following corollary.

COROLLARY 2.4. *[Milchtaich 1996; Ackermann et al. 2009] For any matroid congestion game with altruists and convex delay functions a pure Nash equilibrium exists and can be computed in polynomial time.*

## 3. SINGLETON CONGESTION GAMES

In the previous section we have seen that there are symmetric singleton congestion games with only pure altruists and egoists with and without pure Nash equilibria. For this class of games we can decide the existence of Nash equilibria in polynomial time. In addition, we can compute a Nash equilibrium with minimum and maximum social cost if they exist.

THEOREM 3.1. *For symmetric singleton games with only pure altruists and egoists there is a polynomial time algorithm to decide if a pure Nash equilibrium exists and to compute the best and the worst Nash equilibrium.*

PROOF. We first tackle the existence problem and present an approach similar to [Ieong et al. 2005] based on dynamic programming. The main idea is to reduce the Nash equilibrium property to a constant number of constraints on the congestion values of the machines. These constraints concern the maximum delay of any machine and the minimum delay that an additional altruist or egoist would experience if he arrives on any machine. For each of the polynomially many combinations of these values, we make use of symmetry and implicitly enumerate all allocations that can be a Nash equilibrium. Furthermore, using the social cost to guide the recursion, we are able to find the best and worst Nash equilibria.

Suppose we are given a game $G$ with the set $N_0$ of $n_0$ egoists and the set $N_1$ of $n_1 = n - n_0$ altruists. For a state $S$ consider the set of resources $E_0 = \bigcup_{i \in N_0} S_i$ on which at least one egoist is located. The maximum delay of any resource on which an egoist is located is denoted $d_0^{max} = \max_{e \in E_0} d_e(n_e)$ and minimum delay of any resource if an additional agent is added $d_0^{min+} = \min_{e \in E} d_e(n_e + 1)$. Similarly, consider the set of resources $E_1 = \bigcup_{i \in N_1} S_i$. The maximum altruistic delay of any resource, on which an altruist is located, is denoted $d_1^{max} = \max_{e \in E_1} d_e'(n_e)$ and the minimum altruistic delay

of any resource $d_1^{min+} = \min_{e \in E} d'_e(n_e + 1)$. A state is a Nash equilibrium if and only if

$$d_0^{max} \leq d_0^{min+} \quad \text{and} \quad d_1^{max} \leq d_1^{min+}. \tag{1}$$

This condition yields a separation property. Consider a Nash equilibrium, in which $n_{E',0}$ egoists and $n_{E',1}$ altruists are located on a subset $E' \subset E$ of resources. The Nash equilibrium respects the inequalities above for certain values $d_{0/1}^{max}$ and $d_{0/1}^{min+}$. Note that it is possible to completely change the assignment of agents in $E'$. If the new assignment respects the inequalities for the same values, it can be combined with the assignment on $E - E'$ and again a Nash equilibrium evolves.

This property suggests the following approach to search for an equilibrium. Suppose the values for $d_{0/1}^{max}$ and $d_{0/1}^{min+}$ are given. Our algorithm adds resources $e$ one by one and tests the possible numbers of egoists and altruists that can be assigned to $e$. Suppose we have processed the resources from a subset $E'$ and have found the numbers of altruists and egoists, for which there is an assignment to resources $E'$ such that there is no violation of equations (1) for the given delay values. In this case, we know the feasible numbers of altruists and egoists that are left to be assigned to the remaining resources. Suppose we have marked these combinations of remaining agents in a boolean matrix $R$ of size $(n_0+1) \times (n_1+1)$. Here $r_{ij} = 1$ if and only if there is a feasible assignment of $n_0 - i$ egoists and $n_1 - j$ altruists to $E'$. For the new resource $e$ we now test all combinations $(n_{e,0}, n_{e,1})$ of altruists and egoists that can be allocated to $e$ such that the equations (1) remain fulfilled. We then compile a new matrix $R'$ of the feasible combinations of remaining agents for the remaining resources $E - E' - \{e\}$. In particular, for each tuple $(n_{e,0}, n_{e,1})$ and each positive entry $r_{ij}$ of $R$ we check if $i - n_{e,0} \geq 0$ and $j - n_{e,1} \geq 0$. If this holds, we set the entry of $R'$ with index $(i - n_{e,0}, j - n_{e,1})$ to 1. If $e$ is the last resource to be processed, we check if the resulting matrix $R'$ has a positive entry $r'_{0,0} = 1$. In this case a Nash equilibrium exists for the given values of $d_{\{0,1\}}^{max}$ and $d_{\{0,1\}}^{min+}$, otherwise it does not exist. Due to the separation property mentioned above, this approach succeeds to implicitly test all allocations that fulfill equations (1) for the given values. Finally, note that there are only at most $O(n_0^2 n_1^2 m^4)$ possible values for which we must run the algorithm, as there are at most $(n_0 + 1)m$ possible values for each $d_0^{max}$ and $d_0^{min+}$ and at most $(n_1 + 1)m$ values for each $d_1^{max}$ and $d_1^{min+1}$.

The separation property mentioned above also applies to the best or worst Nash equilibrium. In particular, consider the best Nash equilibrium $S$ that respects (1) for some fixed values $d_{\{0,1\}}^{max}$ and $d_{\{0,1\}}^{min+}$. Consider any subset of resources $E'$ with a number $n_{E',0}$ and $n_{E',1}$ of egoists and altruists, respectively. $S$ is the cheapest Nash equilibrium that respects (1) for the given values if and only if the assignment of $S$ in $E'$ is the cheapest assignment with $n_{E',0}$ egoists and $n_{E',1}$ altruists that respects (1) for the values. Thus, we can adjust our approach as follows. For a set $E'$ of processed resources, instead of simply noting in $r_{ij}$ that there is a feasible assignment to $E'$ that leaves $i$ egoists and $j$ altruists, we can remember the social cost of the cheapest of such assignments. Thus, the matrix $R$ is then a matrix of positive entries, for which we use a prohibitively large cost to identify infeasible combinations. When we compile a new matrix $R'$ after testing all feasible assignments to a new resource $e$, we can denote in each entry the minimum cost that can be obtained for the respective combination. A similar argument works for computing the worst Nash equilibrium. This decides the existence question and finds the cost values of best and worst Nash equilibria. By tracing back the steps of the algorithm we can also discover the strategy choices of agents. □

Note that the previous proof can be extended to a constant number $k$ of different altruism levels. In this more general scenario we choose the delay parameters for each level of altruists. For each resource $e$ we then test all possible combinations of agents from the different levels that we can allocate to a resource $e$ and satisfy all bounds. The matrix $R$ changes in dimension to $(n_{\beta_1}+1)\times\ldots\times(n_{\beta_k}+1)$ to account for all feasible combinations of remaining agents. Finally, we need to test all combinations of delay bounds. However, if $k$ is constant, all these operations can be done in polynomial time.

COROLLARY 3.2. *For symmetric singleton games with altruists and a constant number of different altruism levels, there is a polynomial time algorithm to decide if a pure Nash equilibrium exists and to compute the best and the worst Nash equilibrium.*

As a byproduct, our approach also allows us to compute a social optimum state in polynomial time. We simply assume all agents to be pure altruists and compute the best Nash equilibrium.

COROLLARY 3.3. *For symmetric singleton congestion games a social optimum state can be obtained in polynomial time.*

In case of asymmetric games, however, deciding the existence of pure Nash equilibria becomes significantly harder.

THEOREM 3.4. *It is* NP*-complete to decide if a singleton congestion game with only pure altruists and egoists has a pure Nash equilibrium if $G$ is asymmetric and has concave delay functions.*

PROOF. Membership in NP is obvious. To show completeness, we reduce from 3SAT. Given a formula $\varphi$, we construct a congestion game $G_\varphi$ that has a pure Nash equilibrium if and only if $\varphi$ is satisfiable. Let $x_1,\ldots,x_n$ denote the variables and $c_1,\ldots,c_m$ the clauses of a formula $\varphi$. Without loss of generality [Tovey 1984], we assume each variable appears at most twice positively and at most twice negatively.

For each variable $x_i$ there is a selfish agent $X_i$ that chooses one of the resources $e^1_{x_i}$, $e^0_{x_i}$, or $e_0$. The resources $e^1_{x_i}$ and $e^0_{x_i}$ have the delay function $9x$ and resource $e_0$ has the delay function $7x+3$. For each clause $c_j$, there is a selfish agent $C_j$ who can choose one of the following three resources. For every positive literal $x_i$ in $c_j$ he may choose $e^0_{x_i}$. For every negated literal $\bar{x_i}$ in $c_j$ he may choose $e^1_{x_i}$. Note that there is a stable configuration with no variable agent on $e_0$ if and only if there is a satisfiable assignment for $\varphi$. Additionally, there are three selfish agents $u_1$, $u_2$, and $u_3$ who can choose $e_1$ or $e_2$. Each of the resources $e_1$ and $e_2$ has delay $4$ if used by one agent, delay $8$ if used by two agents and delay $9$ otherwise. The only pure altruist $u_0$ chooses between $e_1$, $e_2$, and $e_0$. Note that the altruist chooses $e_1$, $e_2$ if one of the variable agents is on $e_0$.

If $\varphi$ is satisfiable by a bitvector $(x_1^*,\ldots,x_n^*)$, a stable solution for $G_\varphi$ can be obtained by placing each variable agent $x_i$ on $e^{x_i^*}_{x_i}$. Since $(x_1^*,\ldots,x_n^*)$ satisfies $\varphi$ there is one resource for each clause agent that is not used by a variable agent. Thus, we can place each clause agent on this resource. He then shares this resource with at most one other clause agent due to our assumption on variable appearances in the 3SAT instance. Let the altruist $u_0$ use $e_0$, $u_1$ and $u_2$ choose $e_1$, and $u_3$ choose $e_2$. It is straightforward to verify that this is a Nash equilibrium.

If $\varphi$ is unsatisfiable, there is no stable solution. To prove this it suffices to show that one of the variable agents prefers $e_0$. In that case the altruist never chooses $e_0$ and the agents $u_0,\ldots,u_3$ play the subgame of Example 2.2. For the purpose of contradiction assume that $\varphi$ is not satisfiable but there is a stable solution in which no variable wants to choose $e_0$. This implies that there is no other agent, i.e., a clause agent, on

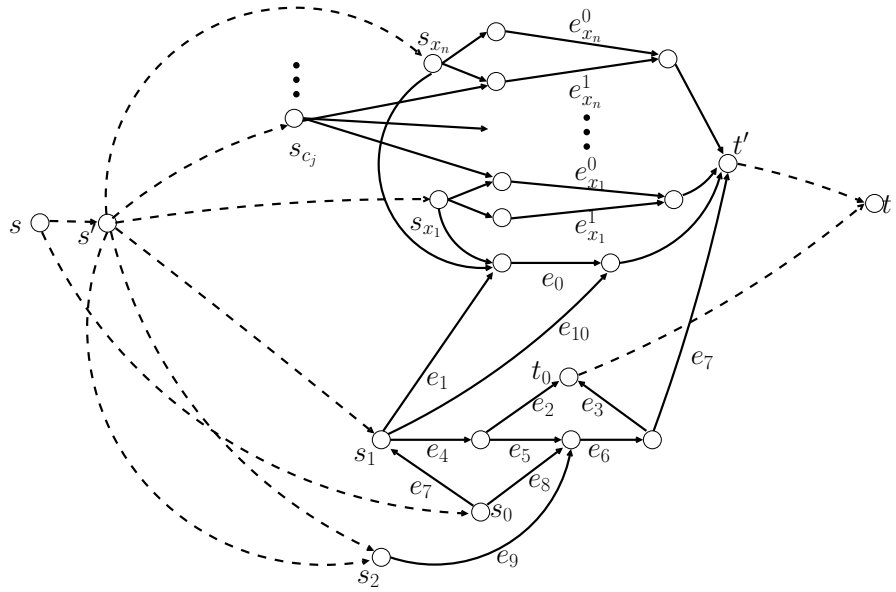


Fig. 1: Structure of the network of $G_\Gamma$ (solid edges only) and $G'_\Gamma$.

a resource that is used by a variable agent. However, if all clause agents are on a resource without a variable agent we can derive a corresponding bit assignment which, by construction, satisfies $\varphi$.

Therefore, $G_\varphi$ has a stable solution if and only if $\varphi$ is satisfiable. □

## 4. GENERAL GAMES

### 4.1. Congestion Games with Convex Delays

For any singleton game $G$ with altruists and convex delay functions, Corollary 2.4 implies that a pure Nash equilibrium always exists. For more general network structures, we show that convexity of delay functions is not sufficient. In particular, this holds even for games with only pure altruists and egoists in the case in which almost all delay functions are linear of the form $d_e(x) = a_e x$, except for two edges, which have quadratic delay functions $d_e(x) = a_e x^2$. For simplicity, we use some edges with constant delay $b_e$. We can replace these edges by sufficiently many parallel edges with delay $b_e x$. This transformation is of polynomial size and yields an equivalent game with only linear and quadratic delay functions.

THEOREM 4.1. *It is* NP*-hard to decide if a symmetric network congestion game with only pure altruists and egoists and quadratic delay functions has a pure Nash equilibrium.*

PROOF. We first reduce from 3SAT to asymmetric congestion games. Again, we assume each variable appears at most twice positively and at most twice negatively. In a second step, we show that the resulting congestion games can be turned into symmetric games while preserving all necessary properties.

Our reduction is similar to the construction that we used in the proof of Theorem 3.4. The structure of the resulting network congestion game $G_\Gamma$ is depicted in Figure 1. The delays are given as in Table I.

| Edge | delay function |
|---|---|
| $e_0$ | $7x+3$ |
| $e_1$ | $2$ |
| $e_2$ | $17$ |
| $e_4$ | $2.4x^2$ |
| $e_6$ | $x^2$ |
| $e_{10}$ | $18.5$ |
| $e^0_{x_i}$,$e^1_{x_i}$ | $9x$ |
| $(s', s_{x_i})$ $\forall 1 \leq i \leq n$ | $Mx$ |
| $(s', s_{c_j})$ $\forall 1 \leq j \leq m$ | $Mx$ |
| $(s', s_1)$, $(s', s_2)$, $(s, s')$ | $Mx$ |
| $(s, s_0)$ | $(n+m+5)M$ |
| $(t_0, t)$ | $(n+m+5)M$ |
| $(t', t)$ | $Mx$ |

Table I: Delay functions on the edges of $G_\Gamma$ and $G'_\Gamma$. Edges not listed here have delay $0$.

Each agent $X_i$ chooses one of three paths from his source node $s_{x_i}$ to his target node $t'$ and therefore uses exactly one of the edges $e^0_{x_i}$, $e^1_{x_i}$, or $e_0$. Each clause agent $C_j$ uses a path from $s_{c_j}$ to $t'$ and uses one of the three edges as described in the proof of Theorem 3.4. That is, for each positive literal $x_i$ in $c_j$ he may choose a path that includes the edge $e^0_{x_i}$. For every negated literal $\bar{x}_i$ in $c_j$ he may choose a path that contains the edge $e^1_{x_i}$.

There is a selfish agent $u_1$ that chooses a path from $s_1$ to $t'$ and two selfish agents $u_2$ and $u_3$ that allocate the path from $s_2$ to $t'$. Finally, one altruistic agent $u_0$ chooses a path from $s_0$ to $t_0$. As in the proof of Theorem 3.4, we can conclude that there is a variable agent whose best response includes edge $e_0$ if and only if $\Gamma$ is not satisfiable. If no variable agent is on $e_0$, a Nash equilibrium can be obtained by placing $u_1$ on the path that begins with $(e_1, e_0)$ which incurs delay of $12$. For agent $u_0$ it is optimal to choose the path $(e_8, e_6, e_3)$ which results in social cost of $39$. This is a Nash equilibrium since the alternative paths are more expensive: for $u_1$ the path via $e_{10}$ costs $18.5$ and the path $(e_4, e_5, e_6, e_7)$ costs $18.4$. For $u_0$ choosing the path $(e_7, e_4, e_2)$ would result in cost of $39.4$.

However, if at least one variable agent is on the edge $e_0$, there is no pure Nash equilibrium. First observe that the path starting with $(e_1, e_0)$ would have delay of at least $19$ for $u_1$ and is always more expensive than the path starting with $e_{10}$. It remains to consider the remaining four strategy profiles of $u_0$ and $u_1$. If the altruist $u_0$ is on the path $(e_8, e_6, e_3)$, the best response for $u_1$ is the path $(e_4, e_5, e_6, e_7)$ which has delay of $18.4$. The alternative path via $e_{10}$ has delay of $18.5$. If $u_1$ is the path $(e_4, e_5, e_6, e_7)$, the best response for the altruist $u_0$ is path $(e_7, e_4, e_2)$ which results in social cost of $63.2$. This is less than the social cost of $66.4$ on the alternative path. If $u_0$ is on $(e_7, e_4, e_2)$, the best response for $u_1$ is the path that begins with the edge $e_{10}$ and has delay of $18.5$. The alternative path $(e_4, e_5, e_6, e_7)$ has delay of $18.6$ This, finally, is a state in which the best response for $u_0$ is $(e_8, e_6, e_3)$ with social cost of $45.5$ which is less than $45.9$ on his alternative path. Thus, the constructed network congestion game $G_\Gamma$ has a pure Nash equilibrium if and only if the formula $\Gamma$ is satisfiable.

Now, we turn the asymmetric network congestion game $G_\Gamma$ into a symmetric congestion game $G'_\Gamma$. We add a new source node $s$, a new target node $t$ and a node $s'$ to the network and connect them to $G_\Gamma$ as depicted by the dashed edges in Figure 1. Note that

$M$ is an integer that is larger than the sum of possible delay values in $G_\Gamma$. If all agents play their best responses, then we can observe the following: Each outgoing edge of $s'$ is used by exactly one selfish agent and the altruist chooses a path that begins with the edge $(s, s_0)$. Every best response path of a selfish agent finishes with the edge $(t', t)$. Every best response path of the altruist ends with the edge $(t_0, t)$. Therefore, $G'_\Gamma$ has a pure Nash equilibrium if and only if $G_\Gamma$ has a pure Nash equilibrium. □

### 4.2. A General Condition for Existence and Convergence

Perhaps surprisingly, if *every* delay function is affine $d_e(x) = a_e x + b_e$, then an elegant combination of the Rosenthal potential and the social cost function yields a potential for arbitrary $\beta_i$-altruists. Hence, existence of Nash equilibria and convergence of sequential better-response dynamics is always guaranteed. The proof carefully exploits the structure of altruistic behavior, as for congestion games with general player-specific affine latency functions a potential does not exist [Gairing et al. 2011].

We prove this result in a more general way for games with weighted potential function. Our construction applies to a general class of these games, in which potential and social cost function have a correlated structure. In particular, we consider potential games with $n$ agents. For simplicity we stick to our notation with $\mathcal{S}_i$ for the strategy space of agent $i$ and $S = (S_1, \ldots, S_n)$ as state of the game. In addition, we assume that there is a personal cost $c_i(S)$ for agent $i$, a potential function $\Phi(S)$, and an arbitrary social cost function $c(S)$. When we consider such games with altruists, an altruistic agent with $\beta_i \in [0, 1]$ again optimizes the trade-off $(1 - \beta_i)c_i(S) + \beta_i c(S)$. While in our examples below we continue to choose $c(S) = \sum_i c_i(S)$, we want to highlight that such a choice is not necessary for our general argument.

We consider the introduction of altruists into potential games with a weighted potential function $\Phi$. In a game with a weighted potential function there is a value $y_i > 0$ for each agent $i$ such that

$$c_i(S) - c_i(S'_i, S_{-i}) = y_i \cdot (\Phi(S) - \Phi(S'_i, S_{-i})) \ ,$$

for every state $S$ and every strategy $S'_i \in \mathcal{S}_i$. In this section we consider games with weighted potential functions of the form

$$\Phi(S) = a \cdot c(S) + \sum_{i=1}^{n} h_i(S_i) \ , \tag{2}$$

where $a > 0$ is a constant and $h_i$ is an arbitrary function that does not depend on any strategy choice $S_j$ by any agent $j \neq i$. While we can always assume $a = 1$ by division of the functions $h_i$ and the values $y_i$, we will allow general values of $a$ for simplicity. The next theorem shows that when we introduce altruists into such games, existence of Nash equilibria and convergence of sequential better-response dynamics is always guaranteed.

THEOREM 4.2. *For every game with weighted potential function, in which social cost and potential function satisfy Equation* (2)*, any corresponding game with altruists is a potential game with weighted potential function*

$$\Psi(S) = c(S) + \sum_{i=1}^{n} \frac{y_i \cdot (1 - \beta_i)}{\beta_i + a \cdot y_i \cdot (1 - \beta_i)} \cdot h_i(S_i) \ ,$$

*where changes in individual cost of agent $i$ and $\Psi$ scale with factor $(\beta_i + a \cdot y_i \cdot (1 - \beta_i))$.*

PROOF. We show that we can transform the potential $\Phi$ for the original game into a potential $\Psi$ for the game with altruists. Consider a state $S$ and an improving strategy

change of an agent $i$ from $S_i$ to $S_i'$ resulting in a strategy profile $S'$. We show that $\Psi$ strictly decreases. For the sake of clarity and brevity we set $\Delta_\Phi = \Phi(S) - \Phi(S')$, $\Delta_C = c(S) - c(S')$, and $\Delta_h = h_i(S_i) - h_i(S_i')$. Note that an improving strategy change for agent $i$ requires

$$\begin{aligned} & (1-\beta_i)(c_i(S) - c_i(S')) + \beta_i(c(S) - c(S')) \\ =\ & y_i(1-\beta_i)\Delta_\Phi + \beta_i\Delta_C \\ >\ & 0\ . \end{aligned}$$

This yields

$$\begin{aligned} \Psi(S) - \Psi(S') &= \left(\Delta_C + \frac{y_i \cdot (1-\beta_i)}{\beta_i + a \cdot y_i \cdot (1-\beta_i)} \cdot \Delta_h\right) \\ &= \frac{1}{\beta_i + a \cdot y_i \cdot (1-\beta_i)} \cdot ((\beta_i + a \cdot y_i \cdot (1-\beta_i))\Delta_C + y_i(1-\beta_i)\Delta_h) \\ &= \frac{1}{\beta_i + a \cdot y_i \cdot (1-\beta_i)} \cdot (a \cdot y_i(1-\beta_i)\Delta_C + y_i(1-\beta_i)\Delta_h + \beta_i\Delta_C) \\ &= \frac{1}{\beta_i + a \cdot y_i \cdot (1-\beta_i)} \cdot (y_i(1-\beta_i)\Delta_\Phi + \beta_i\Delta_C) \\ &> 0\ , \end{aligned}$$

which proves the theorem. □

We proceed to discuss a number of example games from different domains with potential functions given by (2).

*4.2.1. Congestion Games with Affine Delays.* First, let us consider the case of congestion games with altruists and affine delay functions. For simplicity of presentation, we consider linear delays $d_e(x) = a_e x$ without offset $b_e$. This is not a restriction, because (as noted above) we can turn each game with affine delays into an isomorphic game without offsets. If resource $e$ has offset $b_e$, for each player $i$ we introduce a new resource $e_i$ with linear delay $d_{e_i}(x) = b_e x$. Resource $e_i$ is contained only in the strategies of player $i$ that contain $e$. This yields an isomorphic game, in which all individual costs for all states are preserved.

If all delays are linear without offset, then $c(S) = \sum_{e \in E} a_e n_e^2$, whereas the Rosenthal potential reads $\Phi(S) = \frac{1}{2}\sum_{e \in E} a_e(n_e(n_e+1))$. We can set $a = 1/2$, $y_i = 1$, and $h_i(S_i) = \sum_{e \in S_i} a_e/2$ for all agents $i$. With Theorem 4.2 we obtain the weighted potential function

$$\Psi(S) = c(S) + \sum_{i=1}^{n} \frac{1-\beta_i}{1+\beta_i} \sum_{e \in S_i} a_e\ .$$

This yields the following corollary.

COROLLARY 4.3. *For any congestion game with altruists and affine delay functions there is always a pure Nash equilibrium and sequential better-response dynamics converges.*

Unfortunately, it follows directly from previous work [Fabrikant et al. 2004] that the number of iterations to reach a Nash equilibrium can be exponential in the size of the instance, and the problem of computing a Nash equilibrium is PLS-hard. For regular congestion games with matroid strategy spaces [Ackermann et al. 2008] Nash dynamics converge in polynomial time. It is an interesting open problem if a similar result holds for games with altruists.

| | 1 | 2 |
|---|---|---|
| 1 | a,a | d,c |
| 2 | c,d | b,b |

| $\Phi$ | 1 | 2 |
|---|---|---|
| 1 | a-c | 0 |
| 2 | 0 | b-d |

Fig. 2: Payoffs and potential functions for symmetric $2 \times 2$ games.

Note that for convex delay functions the construction of Theorem 4.2 breaks. The reason is that the theorem needs social cost and potential to be related via terms $h_i(S_i)$, which are specific to strategy $S_i$ and independent of $S_{-i}$. For instance, suppose player arrives on a resource with convex delay function. The relation between change in social cost and potential now depends crucially on the number of players on this resource and cannot be expressed via the independent terms $a_e/2$ as for linear delays. As a result, pure Nash equilibria may not exist as observed in Theorem 4.1 above.

*4.2.2. Local Interaction Games.* Another class of games to which the argument can be applied are local interaction games. In a local interaction game [Morris 2000] each player is a node in a graph. He plays a separate symmetric $2 \times 2$ coordination game with every neighbor. However, he can pick his strategy only once and has to apply it in every game he is involved in. Local interaction games are central in the study of cascading behavior and the diffusion of trends in networks (see, e.g., [Kleinberg 2007] for an exposition). Network interaction games [Hoefer and Suri 2012] generalize the construction to games with a possibly different arbitrary symmetric $2 \times 2$ game on each edge. A symmetric $2 \times 2$ game is a 2-player game with $\mathcal{S}_1 = \mathcal{S}_2 = \{1, 2\}$ and arbitrary payoffs given as in Figure 2. Note that these games have an exact potential function, which is displayed in Figure 2. For the social cost function we again assume $c(S) = c_1(S) + c_s(S)$. We set $a = 1/2$, $y_1 = y_2 = 1$. With $h(1) = -c/2$, $h(2) = -d/2$, and $h_1 = h_2 = h$, the function

$$\Psi(S) = c(S) + \sum_{i=1,2} \frac{2(1-\beta_i)}{1+\beta_i} \cdot h_i(S_i)$$

is a weighted potential function for a $2 \times 2$ game with altruists. This construction directly extends to network interaction games, where the cost of a player is summed over different bilateral games played with different players. Thus, Theorem 4.2 yields the following corollary.

COROLLARY 4.4. *For network interaction games with altruists there is always a pure Nash equilibrium and every sequential better-response dynamics converges.*

This class of games generalizes local optimization of weighted MaxCut with the FLIP neighborhood, which is defined by moves of single vertices from one partition to the other. To embed such an instance of weighted MaxCut into our games, every node becomes a player with two strategies representing the two partitions of the cut. Each weighted edge is turned into a bilateral game, where the cost to both players is the edge weight if they pick the same strategy and 0 for both if they pick different strategies. It is straightforward to observe that social cost is twice the cost of edges included in the partitions. Thus, minimizing social cost is equivalent to maximizing the cut weight. A player unilaterally decreases his cost if and only if he decreases his incident edge weight within his partition. Thus, every unilateral deviation is equivalent to a move in the FLIP neighborhood, and it is profitable for the player if and only if it leads to

an increase of cut weight. Thus, every better-response dynamics in this game can be seen as a run of some local search algorithm in the MaxCut instance. This shows that the number of iterations to reach a pure Nash equilibrium can be exponential, and the problem of computing a pure Nash equilibrium is PLS-complete.

The potential function of this game is strikingly similar to the one for congestion games with affine delay functions. This is no coincidence, as local interaction games can be reformulated as congestion games with linear delays. It obviously suffices to model the bilateral $2 \times 2$ games separately. For such a game we introduce four classes of resources with delays $d_{e_1}(x) = (\mathtt{a} - \mathtt{d})x$, $d_{e_{2,i}}(x) = (2\mathtt{d} - \mathtt{a})x$, $d_{e_3}(x) = (\mathtt{b} - \mathtt{c})x$, and $d_{e_{4,i}}(x) = (2\mathtt{c} - \mathtt{b})x$, for $i = 1, 2$. Strategy 1 for agent $i$ contains resources $e_1$ and $e_{2,i}$, strategy 2 contains $e_3$ and $e_{4,i}$. It is straightforward to verify that in every state this yields exactly the same payoffs for all players as the original game.

*4.2.3. Selfish Scheduling.* Our third example comes from the domain of selfish scheduling with coordination mechanisms [Christodoulou et al. 2009]. In this game each agent is a task and chooses one out of $m$ machines as strategy. Thus, the strategy space of each agent is $\mathcal{S}_i = \{1, \ldots, m\}$ for each $i = 1, \ldots, n$. When agent $i$ picks machine $S_i$, the processing time of his task on $S_i$ is $p_{i,S_i} > 0$. There is a local scheduling policy that all machines use to process the tasks that choose to be on the machine. The personal cost $c_i(S)$ of agent $i$ in a state $S$ is the completion time on the chosen machine $S_i$ given the allocation decisions of all agents.

A simple example for a local policy is the *Shortest-First* policy, in which each machine processes the tasks assigned to it without preemption in order of increasing processing time (using a consistent tie-breaking, e.g., based on agent ID). Given the ordering the first task is processed completely, afterwards the second task is started and processed completely, and so on. Recently, a Time-Sharing policy has been proposed by Cohen et al. [2011], in which each machine processes all tasks assigned to it in parallel using processor sharing. Each task that allocates the machine is initially given an equal share of processing time. When a task finishes, then the machine is divided equally among the remaining tasks. Consider for example 3 tasks with processing times 1, 2 and 3 that choose the same machine. With Shortest-First task 1 finishes at time 1, task 2 at time 3 and task 3 at time 6. With Time-Sharing task 1 finishes at time 3, task 2 at time 5, task 3 at time 6. As the social cost of $S$ under the Shortest-First policy (denoted $c^{SF}(S)$) and Time-Sharing policy (denoted $c^{TS}(S)$) we again consider the sum of all player costs.

Both selfish scheduling games with Shortest-First and with Time-Sharing policies are potential games. While for Shortest-First this is established using a lexicographic potential function [Immorlica et al. 2009], games with Time-Sharing allow an exact potential function. A straightforward observation based on cost differences upon migration of a single agent shows that social cost and potential functions satisfy Equation (2). We can reformulate the potential function $\Phi$ for selfish scheduling with Time-Sharing policy given in [Cohen et al. 2011] to

$$\Phi(S) = c^{SF}(S) = \frac{1}{2}\left(c^{TS}(S) + \sum_{i=1}^{n} p_{i,S_i}\right) \quad .$$

With Theorem 4.2 this yields a potential function $\Psi$ for the game with altruists given by

$$\Psi(S) = c^{TS}(S) + \sum_{i=1}^{n} \frac{1-\beta_i}{1+\beta_i} \cdot p_{i,S_i} \quad .$$

COROLLARY 4.5. *For selfish scheduling with Time-Sharing policy and altruists there is always a pure Nash equilibrium and every sequential better-response dynamics converges.*

The similarity of the potential function to the cases before is apparent, and again the game can be transformed into a congestion game with linear delay functions as follows. We first add an agent- and strategy-specific resource $e_{i,S_i}$ to every strategy $S_i$ of agent $i$ with delay $d_{e_{i,S_i}}(x) = 2p_{i,S_i}x$. This accounts for the additional offset in the potential function and the time in $c_i^{TS}(S)$, which is spent on processing the task of agent $i$. The remaining delay of $c_i^{TS}(S)$ can be accounted towards other agents. In particular, when $p_{ij} \geq p_{i'j}$, then agent $i$ is delayed $p_{i'j}$ units due to simultaneous processing of agent $i'$. Similarly, agent $i'$ is delayed only $p_{i'j}$ time units due to the presence of $i$. This allows to construct a local interaction game, in which agents minimize costs, and each pair of agents $i$ and $i'$ plays a symmetric $m \times m$ game. The costs for state $(j, j)$ in this game are $\min\{p_{ij}, p_{i'j}\}$ for both agents and $j = 1, \ldots, m$. The cost for every other state is 0 for both agents. Note that all these local interaction games can easily be turned into congestion games with linear delays. Finally, combining this with the resources $e_{i,S_i}$ completes the construction.

In contrast to general congestion games with linear delays, we show that computing a Nash equilibrium can be done in polynomial time. This result is established with a centralized algorithm that computes a state minimizing the potential. It is an open problem to characterize the duration of better-response dynamics in this class of games.

THEOREM 4.6. *For selfish scheduling with Time-Sharing policy and altruists a Nash equilibrium can be computed in polynomial time.*

PROOF. We show how to compute a state minimizing the potential $\Psi$ given above. This is obviously a pure Nash equilibrium. $\Psi$ consists of the social cost using the Shortest-First policy and some offset terms. This allows to adjust an efficient algorithm for optimizing Shortest-First schedules to compute an optimum for $\Psi$ as well.

The algorithm to find a social optimum schedule for the Shortest-First policy is based on bipartite matching [Bruno et al. 1974] by setting up a complete bipartite network. In this network, one partition is the set of tasks, and the other partition consists of $nm$ nodes $(j, k)$ for positions $k = 1, \ldots, n$ and machines $j = 1, \ldots, m$. The $k$th-to-last position on machine $j$ induces a cost of $k \cdot p_{ij}$ for task $i$. This cost is attached to the corresponding edge $\{i, (j, k)\}$. Note that this is not the cost that $i$ experiences himself but the delay he causes by his presence to himself and the tasks at later positions on machine $j$. For a perfect matching consider the sum of all edge costs. By changing the order of summation, it is straightforward to observe that the total cost of the matching edges is exactly the social cost of the schedule implied by the matching. Hence, a minimum cost perfect matching implies an optimal assignment. A simple pairwise exchange argument shows that the optimal assignment computed by the matching must use a Shortest-First ordering on every machine.

To use this algorithm to minimize our potential, we have to account for the additional terms $\frac{1-\beta_i}{1+\beta_i} \cdot p_{i,S_i}$ in $\Psi$. We do this as follows. We simply add an offset $\frac{1-\beta_i}{1+\beta_i} \cdot p_{ij}$ to each edge weight between task $i$ and any position on machine $j$. Then, by summing the new weights of edges in a perfect matching, we now sum the original edge weights and offsets. The offsets are added up correctly. By reversing the order of summation for the original weights, we get the social cost of the assignment implied by the matching. Note that for a single agent $i$, the offsets are the same for every edge to any position on machine $j$. Thus, the same exchange argument as above shows that the

optimal assignment computed by the matching must use a Shortest-First ordering on every machine. Consequently, this implies that finding a minimum cost perfect matching yields a Shortest-First assignment that minimizes the potential $\Psi$. Thus, we can efficiently compute a global minimum of $\Psi$, which must be a Nash equilibrium. □

For selfish scheduling with Shortest-First policy and altruists it is possible to show that there is no potential function, even for small games with four players and two identical machines, i.e., $p_{i,1} = p_{i,2} = p_i$ for all agents $i$ [Hoefer and Skopalik 2009b].

### 4.3. Weighted Congestion Games with Affine Delays

*4.3.1. Sum-of-Weighted-Delays Social Cost.* In this section we examine the extension to the case of weighted congestion games with linear delay functions. In particular, we consider arbitrary weighted congestion games with delay functions $d_e(x) = a_e x$. Note that the case of affine delays with offsets $b_e$ can again be handled similarly by introducing a player-specific resource for each player $i$ with linear delay $(b_e/w_i)x$. The case of singleton strategies has received a lot of attention under the name KP-model [Koutsoupias and Papadimitriou 2009]. These games were studied as Makespan policy for selfish scheduling [Immorlica et al. 2009]. Weighted congestion games with linear delays have a weighted potential function [Fotakis et al. 2005] given by

$$\Phi(S) = \sum_{e \in E} a_e n_e^2 + \sum_{i=1}^{n} a_{S_i} w_i^2 = \sum_{i=1}^{n} w_i \cdot c_i(S) + \sum_{i=1}^{n} a_{S_i} w_i^2 \ .$$

For agent $i$ we have $c_i(S) - c_i(S'_i, S_{-i}) = \frac{1}{2w_i}(\Phi(S) - \Phi(S'_i, S_{-i}))$. We first consider the social cost function $c^w(S) = \sum_{i=1}^{n} w_i c_i(S)$, which close to what is known in the scheduling literature as sum of weighted completion times. We can use Theorem 4.2 to obtain a potential function $\Psi$ for the game with altruists given by

$$\Psi(S) = c^w(S) + \sum_{i=1}^{n} \frac{2w_i(1-\beta_i)}{\beta_i + 2w_i(1-\beta_i)} \cdot a_{S_i} w_i^2 \ .$$

COROLLARY 4.7. *For weighted congestion games with affine delays and social cost $c^w(S)$ there is always a pure Nash equilibrium and every sequential better-response dynamics converges.*

For the KP-model, Feldmann et al. [2003] proved that for a population of only egoists better-response dynamics can take $O(2^{\sqrt{n}})$ steps to converge to a Nash equilibrium. However, for identical delay functions there is a scheduling of moves to reach a Nash equilibrium with best-response dynamics in polynomial time. In addition, there are polynomial time algorithms to compute Nash equilibria for asymmetric singleton games with linear delay functions [Feldmann et al. 2003; Gairing et al. 2010].

*4.3.2. Sum-of-Delays Social Cost.* As a second social cost function let us again consider $c(S) = \sum_i c_i(S)$, which has been done previously for this model in [Hoefer and Souza 2010; Berenbrink et al. 2006]. This cost function is close to what is known in the scheduling literature as sum of completion times. Although in this case the relation between potential and social cost function is only slightly different from the condition in Equation (2), our results are mostly negative. We observe that even for identical delays and only one altruist, existence of a pure Nash equilibrium is not guaranteed.

*Example* 4.8. Consider a game with two edges, one pure altruist with $w_1 = 5$, and four egoists with $w_2 = 10$, $w_3 = w_4 = w_5 = 1$. Assume there is a pure Nash equilibrium,

then agent $2$ chooses a different edge than agent 1. The agents 3, 4, and 5 choose a different edge than agent 2. However, agent 1 would choose the machine with only agent 2, which leads to a contradiction. The idea can be adjusted to an arbitrary altruist with $\beta_1 > 0$ by adding sufficiently many agents with small weight. In particular, instead of 3 we add strictly more than $1 + \frac{1.4}{\beta_1}$ many egoists, which all have equally small weight, and for which their total weight adds up to 3. For this game it can be shown that all arguments given above are preserved.

In addition, we can show that it is NP-complete to decide if a pure Nash equilibrium exists. The reduction is from PARTITION.

THEOREM 4.9. *It is* NP*-complete to decide if a weighted congestion game on three identical parallel links with linear delays and one pure altruist has a pure Nash equilibrium.*

PROOF. Membership in NP is obvious. To show completeness, we reduce from PARTITION. An instance $\mathcal{I}$ is given as $(a_1, \ldots, a_n) \in \mathbb{N}^n$ and $\mathcal{I} \in$ PARTITION if and only if $\exists I \subset \{1, \ldots, n\}$ with $\sum_{i \in I} a_i = \sum_{j \in \{1,\ldots,n\} \setminus I} a_j$. We first reduce a given instance $\mathcal{I} = (a_1, \ldots, a_n)$ to an instance $\mathcal{I}' = (a_1, \ldots, a_n, a_{n+1}, \ldots, a_{n+8})$ with $a_{n+1} = \ldots = a_{n+8} = \sum_{i \in \{1,\ldots,n\}} a_i$. Clearly $\mathcal{I} \in$ PARTITION if and only if $\mathcal{I}' \in$ PARTITION.

In a second step we construct a game $G_{\mathcal{I}'}$ that has a Nash equilibrium if and only if $\mathcal{I}' \in$ PARTITION. The game consists of three edges and $n + 8 + 2$ agents. The weight $w_i$ of agent $1 \leq i \leq n + 8$ is $a_i$. Agent $n + 9$ has weight $p_{n+9} = \sum_{1 \leq j \leq n+8} a_j$ and agent $n + 10$ has weight $p_{n+10} = \frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$. All agents are pure egoists except for task $n + 10$ who is a pure altruist.

If $\mathcal{I} \in$ PARTITION, there is an $I \subset \{1, \ldots, n + 8\}$ with $\sum_{i \in I} a_i = \frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$. Assigning all agents $i \in I$ to edge one, all agents $j \in \{1, \ldots, n + 8\} \setminus I$ to edge two, and the remaining agents $n + 9$ and $n + 10$ to edge three is a Nash equilibrium. Note that the first two edges have a total weight of $\frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$ and edge three has a weight of $\frac{3}{2} \sum_{1 \leq j \leq n+8} a_j$. Obviously, no agent from the first two edges has an incentive to change to edge three. Neither has agent $n+9$ an incentive to change to one of the first two edges because his delay would not change. The altruistic agent cannot improve the social cost by changing to one of the first two edges. Note that at least $4$ agents (half of the agents $n, \ldots, n + 8$) are assigned to each of the first two edges. Therefore, when the altruistic agents migrates the social cost increases by at least $4w_{n+10} - (w_{n+10} + w_{n+9}) > 0$.

If $\mathcal{I} \notin$PARTITION, assume for the sake of contradiction that there is a Nash equilibrium. Observe that agent $n + 9$ does not choose the edge that agent $n + 10$ is on. Since there is no $I \subset \{1, \ldots, n + 8\}$ with $\sum_{i \in I} a_i = \frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$, there exists an edge that has congestion of less than $\frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$ (while ignoring agent $n + 9$). On the other hand, each of the agents $1, \ldots, n + 8$ can always choose an edge that has congestion less than $p_{n+9}$. Therefore, in equilibrium they choose the other two edges. Note, that each of these two edges has at least $4$ of the agents $n, \ldots, n + 8$. Finally, the altruistic agent $n + 10$ chooses the edge that only agent $n + 9$ is assigned to (changing to one of the other two edges increases the social cost by at least $4w_{n+10} - (w_{n+10} + w_{n+9}) > 0$). This contradicts the existence of a Nash equilibrium. □

Our proof requires the presence of agents with different altruism levels. It is an interesting open problem if the existence of a potential function can be shown for games with $\beta$-uniform altruists.

## 5. STABILIZATION METHODS

This section treats a scenario in which a central institution can convince selfish agents to adopt socially favorable behavior and act altruistically. For simplicity of presentation we first restrict to games with only pure altruists and egoists. This is closely related to Stackelberg routing [Korilis et al. 1997; Roughgarden 2004; Fotakis 2010] in atomic congestion games, in which an altruistic leader can control the strategy choice of a fixed portion of agents and strives to minimize the overall cost of the resulting Nash equilibrium.

A natural question in our scenario is how many altruists are required to guarantee that there is a Nash equilibrium with a certain cost, e.g. a Nash equilibrium as cheap as a social optimum state. This problem has been considered in non-atomic congestion games in [Kaporis and Spirakis 2009; Sharma and Williamson 2009]. We term this number the *optimal stability threshold*. In a more pessimistic direction it is of interest to determine the minimum number of altruists needed to guarantee that the worst-case Nash equilibrium is optimal. We term this number the *optimal anarchy threshold*. Let us denote by $n_1^+$ and $n_1^-$ the optimal stability and anarchy threshold, respectively. As a consequence from Theorem 3.1 we can compute both numbers for symmetric singleton congestion games in polynomial time. For each number of altruists we check if the best and/or worst Nash equilibrium is as cheap as the social optimum.

COROLLARY 5.1. *For symmetric singleton congestion games with only pure altruists and egoists there is a polynomial time algorithm to compute $n_1^+$ and $n_1^-$.*

Note that the optimal anarchy threshold is not well-defined, because the worst Nash equilibrium might always be suboptimal, even for a population of altruists only. In case of symmetric singleton games and convex delay functions, an easy exchange argument serves to show that in this case any local optimum is also a global optimum. However, for concave delay functions or asymmetric singleton games, a local optimum might still be globally suboptimal.

*Example* 5.2. Consider a symmetric game with two resources, $d_1(1) = 16$, $d_1(2) = 32$, $d_1(3) = 36$, and $d_2(x) = 45$. If all agents allocate resource 1, we get a Nash equilibrium of cost 108. In the optimum two agents allocate resource 2 resulting in a cost of 106. Now consider an asymmetric game with three resources and delay functions $d_1(x) = d_2(x) = 8x$, and $d_3(x) = 4x$. Agent 1 can use resources 1 and 2, agents 2 and 3 can use resources 2 and 3. The state $(2, 3, 3)$ is a Nash equilibrium of cost 32, while the social optimum is a state $(1, 2, 3)$ of cost 20.

Note that for symmetric games, our algorithm is able to detect the cases in which suboptimal local optima exist. In the asymmetric case, however, a similar approach fails, because of the NP-hardness of determining existence of a pure Nash equilibrium. Thus, in the following we concentrate on the optimal stability threshold.

In asymmetric games, it is also required to determine the identity of agents, so here we strive to find a set (denoted $N_e^+$) of minimum cardinality. For an optimal set of congestion values $n_E^* = (n_e^*)_{e \in E}$ we can determine $N_1^+(n_E^*)$ such that there is a Nash equilibrium of the game with congestion values $n_e^*$ for all $e \in E$.

THEOREM 5.3. *For singleton games with only pure altruists and egoists and a social optimal congestion vector $n_E^*$ there is a polynomial time algorithm to compute $N_1^+(n_E^*)$.*

PROOF. Suppose we are given a congestion vector $n_E^*$ that results in minimum social cost. Our aim is to derive a state of the game (i.e., an assignment agents to resources) such that the state is a Nash equilibrium, we have congestion $n_e^*$ on every resource $e$, and a minimal set of altruists are needed. Note that whenever we assign an agent to

become an altruist, we are free to assign him to any resource – if the resulting state has the social optimal congestion vector $n_E^*$, an altruist has no incentive to deviate in this state from any strategy he is playing. For an egoist, we can determine – given that the final state has congestion vector $n_E^*$ – which of the resources would qualify as a best response for him. That is, in the final state he needs to be assigned to a resource $e \in \mathcal{S}_i$ such that $d_e(n_e^*) \leq d_{e'}(n_{e'}^* + 1)$ for every $e' \in \mathcal{S}_i$.

This allows to compute $N_1^+(n_E^*)$ by constructing a weighted bipartite graph as follows. One partition is the set of agents $N$. In the other partition we introduce for each resource $e$ a number of $n_e^*$ vertices. If $e \in \mathcal{S}_i$ we connect agent $i$ to all vertices that were introduced due to $e$. Suppose $e$ represents a best response for $i$ with respect to egoistic delay $d$. Then we assign a weight of 0 to all corresponding edges between $i$ and the vertices of $e$. To all other edges we assign a weight of 1. Note that any feasible allocation of agents to strategies that generates the congestion vector $n_E^*$ is represented by a perfect matching. If we match an agent to a strategy, which is not an egoistic best response, it has to become an altruist and a weight of one is counted towards the weight of the matching. By computing a minimum weight perfect matching [Cook and Rohe 1999], we can identify a minimal set $N_1^+(n_E^*)$ of altruists required to stabilize $n_E^*$. □

Observe that by creating the edges of cost 1 only to strategies which represent best responses with respect to the altruistic delay $d'$, we can compute $N_1^+(n_E)$ for arbitrary congestion vectors $n_E$. In this case, the set might be empty, if, e.g., the congestion vector corresponds to a very expensive state and can never be generated by a Nash equilibrium for any distribution of altruists. This case, however, can be recognized by the absence of a perfect matching in the bipartite graph.

This approach turns out to be applicable to an even more general natural scenario. Suppose each agent $i$ has a *stability cost* $c_{ie}$ for each strategy $e \in \mathcal{S}_i$. This cost yields the disutility for being forced to play a certain strategy given a congestion vector $n_E$. In this scenario we slightly change $N_1^*(n_E)$ to the set agents of minimal stability cost. Still, we can compute this set by a minimum weight perfect matching if we set the weights to $c_{ie}$ for all edges connecting $i$ to vertices of $e$. The stability cost allows for general preferences exceeding categories like altruists and egoists.

COROLLARY 5.4. *For singleton games and a congestion vector $n_E$ there is a polynomial time algorithm to compute $N_1^+(n_E)$ with minimal stability cost.*

The underlying problem can be seen as a slot allocation to agents. As the computed allocation has minimal stability cost, it is possible to turn the algorithm into a truthful mechanism using VCG payments (see e.g. [Nisan et al. 2007, chapter 9]). Our final mechanism (1) learns the stability costs from each agent, (2) determines the allocation, and (3) pays appropriate amounts to agents for truthful revelation of cost values and adaptation of allocated strategies. Computing the optimal allocation as well as the payments are bipartite matching problems (computing an optimum solution, computing optimal solutions excluding single agents). All of these problems can be solved in polynomial time, and hence the resulting mechanism can be implemented in polynomial time.

COROLLARY 5.5. *For singleton games and a congestion vector $n_E$ there is a truthful VCG-mechanism to compute $N_1^+(n_E)$ in polynomial time.*

These general results are restricted to the case of singleton games. For more general games we show that it is NP-hard to decide if there is a Nash equilibrium as cheap as the social optimum. Our next theorem establishes this even for symmetric network congestion games with linear delays, in which an arbitrary Nash equilibrium and a social optimum state can be computed in polynomial time [Fabrikant et al. 2004].

Furthermore, the result requires only a series-parallel network. Thus, even in this restricted case it is NP-hard to decide if the number $n_1^+$ of pure altruists required is 0 or 1, or equivalently if $N_1^+(n_E^*)$ is empty or not. This directly yields hardness of approximation within any finite factor.

THEOREM 5.6. *For symmetric network congestion games with 3 agents, linear delay functions on series-parallel graphs and optimal congestions $n_E^*$ it is* NP*-hard to decide if there is a pure Nash equilibrium with congestions $n_E^*$.*

PROOF. We reduce from PARTITION. Let an instance be given by positive integers $a_1, \ldots, a_k$ and $a = \sum_{i=1}^{k} a_i$, where $a$ is an even number. Create a network with two nodes and two parallel edges $e_1$ and $e_2$ for each integer $a_i$. The delay $d_{e_1}(x) = 2a_i x$, and $d_{e_2}(x) = a_i x$. All these networks are concatenated sequentially. We denote the first node of this path gadget by $u$ and the last by $v$. In addition, we add one edge $f = (u, v)$ with delay $d_f(x) = \frac{3}{4}ax$. Finally, the game has three egoists, which need to allocate a path from $u$ to $v$.

The unique social optimum is to let one agent use $f$ and the other two agents use two edge-disjoint paths through the path gadget. This yields an optimal social cost of $\frac{15}{4}a$. However, for a Nash equilibrium each path through the gadget must not have more delay than $\frac{3}{2}a$. If the instance of PARTITION is solvable, then the elements assigned to a partition represent the edges of type $e_1$ that an agent allocates in Nash equilibrium. Otherwise, if the instance is not solvable, there is no possibility to partition the path gadget into two edge-disjoint paths of latency at most $\frac{3}{2}a$.

The reduction works for a small constant number of agents but only shows weak NP-hardness. If the number of agents is variable, it is possible to show strong NP-hardness with a similar reduction from 3-PARTITION. □

We remark that the previous theorem contrasts the continuous non-atomic case, in which a minimal fraction of altruistic demand stabilizing an optimum solution can be computed in any symmetric network congestion game [Kaporis and Spirakis 2009].

## 6. CONCLUSIONS

In this paper, we have initiated the study of altruists in atomic congestion games. Our model is similar to the one presented by Chen and Kempe [2008] for nonatomic routing games, however, we observe quite different properties. In the nonatomic case, existence of Nash equilibria for any population of agents is always guaranteed, even if agents are partially spiteful. In contrast, our study answers fundamental questions for existence and convergence in atomic games. For the case of linear latencies, an elegant combination of social cost and the Rosenthal potential proves guaranteed existence and convergence. In addition, this result is based on a more general condition, which is applicable to prove existence and convergence in games with potential function from different domains. In the case of weighted congestion games on parallel links with linear delay functions we have observed that even a slight variation in the interplay of social cost and potential functions can lead to instability and negative results.

There are a number of open problems and research directions that stem from the results in this paper. An interesting open problem is to consider the relations to results on Stackelberg games [Fotakis 2010]. While we have studied existence of pure Nash equilibria and convergence of better-response dynamics in congestion games with sum social cost, the natural open problem is to analyze the model for other prominent social cost functions – with maximum cost of any player being the prime example. In addition, there are several generalizations such as bottleneck congestion games, in which consideration of altruistic behavior might be worthwhile. More generally, it is important to obtain a deeper understanding of more general (potential) games, in which pure

Nash equilibria exist and convergence is guaranteed even for altruistic agents. Finally, the complexity of computing such an equilibrium and the duration of best- and better-response dynamics in classes of games with altruistic agents represent intriguing open problems for further research.

## REFERENCES


ACKERMANN, H., RÖGLIN, H., AND VÖCKING, B. 2008. On the impact of combinatorial structure on congestion games. *J. ACM 55,* 6.

ACKERMANN, H., RÖGLIN, H., AND VÖCKING, B. 2009. Pure Nash equilibria in player-specific and weighted congestion games. *Theoret. Comput. Sci. 410,* 17, 1552–1563.

BABAIOFF, M., KLEINBERG, R., AND PAPADIMITRIOU, C. 2009. Congestion games with malicious players. *Games Econom. Behav. 67,* 1, 22–35.

BERENBRINK, P., GOLDBERG, L. A., GOLDBERG, P., AND MARTIN, R. 2006. Utilitarian resource assignment. *J. Disc. Algorithms 4,* 4, 567–587.

BRUNO, J., COFFMAN JR., E., AND SETHI, R. 1974. Scheduling independent tasks to reduce mean finishing time. *Comm. ACM 17,* 7, 382–387.

CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., AND PAPAIOANNOU, E. 2010. The impact of altruism on the efficiency of atomic congestion games. In *Proc. 5th Symp. Trustworthy Global Computing (TGC)*. 172–188.

CHEN, P.-A., KEIJZER, B. D., KEMPE, D., AND SCHAEFER, G. 2011. On the robust price of anarchy of altruistic games. In *Proc. 7th Intl. Workshop Internet & Network Economics (WINE)*. 383–390.

CHEN, P.-A. AND KEMPE, D. 2008. Altruism, selfishness, and spite in traffic routing. In *Proc. 9th Conf. Electronic Commerce (EC)*. 140–149.

CHRISTODOULOU, G., KOUTSOUPIAS, E., AND NANAVATI, A. 2009. Coordination mechanisms. *Theoret. Comput. Sci. 410,* 36, 3327–3336.

COHEN, J., DÜRR, C., AND KIM, T. N. 2011. Non-clairvoyant scheduling games. *Theory Comput. Syst. 49,* 1, 3–23.

COOK, W. AND ROHE, A. 1999. Computing minimum-weight perfect matchings. *INFORMS J. Comput. 11,* 2, 138–148.

ESHEL, I., SAMUELSON, L., AND SHAKED, A. 1998. Altruists, egoists and hooligans in a local interaction model. *Amer. Econ. Rev. 88,* 1, 157–179.

FABRIKANT, A., PAPADIMITRIOU, C., AND TALWAR, K. 2004. The complexity of pure Nash equilibria. In *Proc. 36th Symp. Theory of Computing (STOC)*. 604–612.

FEHR, E. AND SCHMIDT, K. 1999. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics 114*, 817–868.

FELDMANN, R., GAIRING, M., LÜCKING, T., MONIEN, B., AND RODE, M. 2003. Nashification and the coordination ratio for a selfish routing game. In *Proc. 30th Intl. Coll. Automata, Languages and Programming (ICALP)*. 514–526.

FOTAKIS, D. 2010. Stackelberg strategies for atomic congestion games. *Theory Comput. Syst. 47,* 1, 218–249.

FOTAKIS, D., KONTOGIANNIS, S., AND SPIRAKIS, P. 2005. Selfish unsplittable flows. *Theoret. Comput. Sci. 348,* 2–3, 226–239.

GAIRING, M. 2008. Malicious Bayesian congestion games. In *Proc. 6th Intl. Workshop Approx. and Online Algorithms (WAOA)*. 119–132.

GAIRING, M., LÜCKING, T., MAVRONICOLAS, M., AND MONIEN, B. 2010. Computing Nash equilibria for scheduling on restricted parallel links. *Theory Comput. Syst. 47,* 2, 405–432.

GAIRING, M., MONIEN, B., AND TIEMANN, K. 2011. Routing (un-)splittable flow in games with player-specific linear latency functions. *ACM Trans. Algorithms 7,* 3, 31.

GINTIS, H., BOWLES, S., BOYD, R., AND FEHR, E. 2005. *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press.

HOEFER, M. AND SKOPALIK, A. 2009a. Altruism in atomic congestion games. In *Proc. 17th European Symp. Algorithms (ESA)*. 179–189.

HOEFER, M. AND SKOPALIK, A. 2009b. Stability and convergence in selfish scheduling with altruistic agents. In *Proc. 5th Intl. Workshop Internet & Network Economics (WINE)*. 616–622.

HOEFER, M. AND SOUZA, A. 2010. Tradeoffs and average-case equilibria in selfish routing. *ACM Trans. Comp. Theory 2,* 1.

HOEFER, M. AND SURI, S. 2012. Dynamics in network interaction games. *Distributed Computing 25,* 5, 359–370.

IEONG, S., MCGREW, R., NUDELMAN, E., SHOHAM, Y., AND SUN, Q. 2005. Fast and compact: A simple class of congestion games. In *Proc. 20th Conf. Artificial Intelligence (AAAI)*. 489–494.

IMMORLICA, N., LI, L., MIRROKNI, V., AND SCHULZ, A. 2009. Coordination mechanisms for selfish scheduling. *Theoret. Comput. Sci. 410,* 17, 1589–1598.

KAPORIS, A. AND SPIRAKIS, P. 2008. Stackelberg games: The price of optimum. In *Encyclopedia of Algorithms*. Springer Verlag.

KAPORIS, A. AND SPIRAKIS, P. 2009. The price of optimum in Stackelberg games on arbitrary single commodity networks and latency functions. *Theoret. Comput. Sci. 410,* 8–10, 745–755.

KARAKOSTAS, G. AND VIGLAS, A. 2007. Equilibria for networks with malicious users. *Math. Prog. 110,* 3, 591–613.

KLEINBERG, J. 2007. Cascading behavior in netwoks: Algorithmic and economic issues. See Nisan et al. [2007], Chapter 24.

KORILIS, Y., LAZAR, A., AND ORDA, A. 1997. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Trans. Netw. 5,* 1, 161–173.

KOUTSOUPIAS, E. AND PAPADIMITRIOU, C. 2009. Worst-case equilibria. *Computer Science Review 3,* 2, 65–69.

LEDYARD, J. 1997. Public goods: A survey of experimental resesarch. In *Handbook of Experimental Economics*, J. Kagel and A. Roth, Eds. Princeton University Press, 111–194.

LEVINE, D. 1998. Modeling altruism and spitefulness in experiments. *Rev. Econom. Dynamics 1*, 593–622.

MILCHTAICH, I. 1996. Congestion games with player-specific payoff functions. *Games Econom. Behav. 13,* 1, 111–124.

MORRIS, S. 2000. Contagion. *Rev. Econ. Stud. 67,* 1, 57–78.

MOSCIBRODA, T., SCHMID, S., AND WATTENHOFER, R. 2009. The price of malice: A game-theoretic framework for malicious behavior in distributed systems. *Internet Math. 6,* 2, 125–155.

NISAN, N., TARDOS, É., ROUGHGARDEN, T., AND VAZIRANI, V., Eds. 2007. *Algorithmic Game Theory*. Cambridge University Press.

ROSENTHAL, R. 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory 2*, 65–67.

ROUGHGARDEN, T. 2004. Stackelberg scheduling strategies. *SIAM J. Comput. 33,* 2, 332–350.

SHARMA, Y. AND WILLIAMSON, D. 2009. Stackelberg thresholds in network routing games or the value of altruism. *Games Econom. Behav. 67,* 1, 174–190.

TOVEY, C. 1984. A simplified NP-complete satisfiability problem. *Disc. Appl. Math. 8*, 85–89.

Received April 2012; revised August 2012; accepted October 2012