Convergence Issues in Competitive Games

Vahab S. Mirrokni[1] and Adrian Vetta[2]

[1] Massachusetts Institute of Technology: mirrokni@theory.lcs.mit.edu
[2] McGill University: vetta@math.mcgill.ca

Abstract. We study the speed of convergence to approximate solutions in iterative competitive games. We also investigate the value of Nash equilibria as a measure of the cost of the lack of coordination in such games. Our basic model uses the underlying best response graph induced by the selfish behavior of the players. In this model, we study the value of the social function after multiple rounds of best response behavior by the players. This work therefore deviates from other attempts to study the outcome of selfish behavior of players in non-cooperative games in that we dispense with the insistence upon only evaluating Nash equilibria. A detailed theoretical and practical justification for this approach is presented. We consider non-cooperative games with a submodular social utility function; in particular, we focus upon the class of valid-utility games introduced in [13]. Special cases include basic-utility games and market sharing games which we examine in depth. On the positive side we show that for basic-utility games we obtain extremely quick convergence. After just one round of iterative selfish behavior we are guaranteed to obtain a solution with social value at least $\frac{1}{3}$ that of optimal. For n-player valid-utility games, in general, after one round we obtain a $\frac{1}{2n}$-approximate solution. For market sharing games we prove that one round of selfish response behavior of players gives $\Omega(\frac{1}{\ln n})$-approximate solutions and this bound is almost tight. On the negative side we present an example to show that even in games in which every Nash equilibrium has high social value (at least half of optimal), iterative selfish behavior may "converge" to a set of extremely poor solutions (each being at least a factor n from optimal). In such games Nash equilibria may severely underestimate the cost of the lack of coordination in a game, and we discuss the implications of this.

1 Introduction

Traditionally, research in operation research has focused upon finding a global optimum. Computer scientists have also long studied the effects of lack of different resources, mainly the lack of computational resources, in optimization. Recently, the *lack of coordination* inherent in many problems has become an important issue in computer science. A natural response to this has been to analyze Nash equilibria in these games. Of particular interest is the *price of anarchy* in a game [8]; this is the worst case ratio between an optimal social solution and a Nash equilibrium. Clearly, a low price of anarchy may indicate that a system has no need for a single regulatory authority. Conversely, a high price of anarchy is indicative of a poorly functioning system in need of some regulation.

In this paper we move away from the use of Nash equilibria as the solution concept in a game. There are several reasons for this. The first reason relates to use of non-randomized (pure) and randomized (mixed) strategies. Often pure Nash equilibria may not exist, yet in many games the use of a randomized (mixed) strategy is unrealistic. This necessitates the need for an alternative in evaluating such games.

Secondly, Nash equilibria represent stable points in a system. Therefore (even if pure Nash equilibria exist), they are a more acceptable solution concept if it is likely that the system does converge to such stable points. In particular, the use of Nash equilibria seems more valid in games in which Nash equilibria arise when players iteratively engage in selfish behavior. The time it takes for *convergence* to Nash equilibria, however, may be extremely long. So, from a practical viewpoint, it is important to evaluate the speed or rate of convergence. Moreover, in many games it is not the case that repeated selfish behavior always leads to Nash equilibria. In these games, it seems that another measure of the cost of the lack of coordination would be useful.

As is clear, these issues are particularly important in games in which the use of pure strategies and repeated moves are the norm, for example, auctions. We remark that for most practical games these properties are the rule rather than the exception (this observation motivates much of the work in this paper). For these games, then, it is not sufficient to just study the value of the social function at Nash equilibria. Instead, we must also investigate the speed of convergence (or non-convergence) to an equilibrium. Towards this goal, we will not restrict our attention to Nash equilibria but rather prove that after some number of improvements or best responses the value of the social function is within a factor of the optimal social value. We tackle this by modeling the behavior of players using the underlying best response graph on the set of strategy states. We consider (best response) paths in this graph and evaluate the social function at states along these paths. The rate of convergence to high quality solutions (or Nash equilibria) can then be measured by the length of the path. As mentioned, it may the case that there is no such convergence. In fact, in Section 4.2, it is shown that instead we have the possibility of "convergence" to non-Nash equilibria with a bad social value. Clearly such a possibility has serious implications for the study of stable solutions in games.

An overview of the paper is as follows. In section 2, we describe the problem formulations and model. In section 3, we discuss other work and their relation to this paper. In section 4, we give results for valid-utility and basic-utility games. We prove that in valid-utility games we obtain a $\frac{1}{2n}$-approximate solution if each player sequentially makes one best response move. For basic-utility games we obtain a $\frac{1}{3}$-approximate solution in general, and a $\frac{1}{2}$-approximate solution if each player initially used a null strategy. We then present a valid-utility game in which every Nash equilibria is at least half-optimal and, yet, iterative selfish behavior may lead to only $O(\frac{1}{n})$-approximate solutions. In section 5, we examine market sharing games and show that we obtain $\Omega(\frac{1}{\ln n})$-approximate solutions after one best response move each. Finally, in section 6, we discuss other classes of games and present some open questions.

2 Preliminaries

In this section, we define necessary game theoretic notations to formally describe the classes of games that we study in the next sections. The game is defined as the tuple $(U, \{S_j\}, \{\alpha_j()\})$. Here U is the set of players or agents. Associated with each player j is a disjoint groundset V_i, and S_j is a collection of subsets of V_j. The elements in the a groundset correspond to acts a player may make, and hence the subsets correspond to strategies. We denote player j's strategy by $s_j \in S_j$. Finally, $\alpha_j : \Pi_j S_j \to \mathbb{R}$ is the private payoff or utility function for agent j, given the set of actions of all the players. In a non-cooperative game, we assume that each selfish agent wishes to maximize its own payoff.

Definition 1. *A function $f : 2^V \to \mathbb{R}$ is a set function on the groundset V. A set function $f : 2^V \to \mathbb{R}$ is submodular if for any two sets $A, B \subseteq V$, we have $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. The function is non-decreasing if $f(X) \leq f(Y)$ for any $X \subseteq Y \subseteq V$.*

For each game we will have a social objective function $\gamma : \Pi_j S_j \to \mathbb{R}$. (We remark that γ can be viewed as a set function on the groundset $\cup V_i$.) Our goal will be to analyze the social value of solutions produced the selfish behavior of the agents. Specifically, we will focus upon the class of games called *valid-utility games.*

Definition 2. *Let $\mathcal{G}(U, \{S_j\}, \{\alpha_j\})$ be a non-cooperative game with social function γ. $\mathcal{G}$ is a valid-utility game if it satisfies the properties:*

- *γ is a submodular set function.*
- *The payoff of a player is at least equal to the difference in the social function when the player participates versus when it does not.*
- *The sum of the utility or payoff functions for any set of strategies should be less than or equal to the social function.*

This framework encompasses a wide range of games in facility location, traffic routing and auctions [13]. Here, as our main application, we consider the market sharing game which is a special case of valid-utility games (and also congestion games). We define this game formally in Section 5.

2.1 Best Response Paths

We model the selfish behavior of players using an underlying *state graph.* Each vertex in the graph represents a strategy state $S = (s_1, s_2, \ldots, s_n)$. The arcs in the graph corresponds to best response moves by the players. Formally, we have

Definition 3. *The state graph, $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, is a directed graph. Each vertex in $\mathcal{V}$ corresponds to a strategy state. There is an arc from state S to state S' with label j if the only difference between $\mathcal{S}$ and $\mathcal{S}'$ is only in the strategy of player j; and player j plays his best response in strategy state S to go to S'.*

Observe that the state graph may contain loops. A best response path is a directed path in the state graph. We say that a player i plays in the best response path $\mathcal{P}$, if at least one of the edges of $\mathcal{P}$ is labelled i. Assuming that players optimize their best response function sequentially (and not in parallel), we can evaluate the social value of states on a best response path in the state graph. In particular, given a best response path starting from an arbitrary state, we will be most interested in the social value of the the last state on the path. Notice that if we do not allow every player to make a best response on a path $\mathcal{P}$ then we may not be able to bound the social value of a state with respect to the optimal solution. This follows from the fact that the actions of a single player may be very important for producing solutions of high social value. Hence, we consider the following models:

One-round path: Consider an arbitrary ordering of players $i_1, \ldots, i_n$. Path $\mathcal{P}$ is a *one-round path* if it starts from an arbitrary state and edges of P are labelled $i_1, i_2, \ldots, i_n$ in this order.

Covering path: A best response path $\mathcal{P}$ is a *covering path* if each player plays at least once on the path.

k-Covering path: A best response path $\mathcal{P}$ is a *k-covering path* if there are k covering paths $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_k$ such that $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_k)$.

Observe that a one-round path is a covering path. Note that in the one-round path we let each player play his best response exactly one time, but in a covering path we let each player play at least one time. Both of these models have justifications in extensive games with complete information. In these games, the action of each player is observed by all the other players. As stated, for a non-cooperative game $\mathcal{G}$ with a social function γ, we are interested in the social value of states (especially the final state) along one-round, covering, and k-covering paths.

A Simple Example. Here, we consider covering paths in a basic load balancing game; Even-Dar et. al. [2] considered the speed of convergence to Nash equilibria in these games. There are n jobs that can be scheduled on m machines. It takes p_j units of time for job j to run on any of the machines. The social objective function is the maximum makespan over all machines. The private payoff of a job, however, is the inverse of the makespan of the machine that the job is scheduled on. Thus each job wants to be scheduled on a machine with as small a makespan as possible. It is easy to verify that the price of anarchy in this game is at most 2. It is also known that this game has pure Nash equilibria and the length of any best-response path in this game is at most n^2 [1]. In addition, from any state there is a path of length at most n to some pure Nash equilibrium [12]. It may, however, take much more than n steps to converge to a pure Nash equilibrium. Hence, our goal here is to show that the social value of any state at the end of a covering path is within a factor 2 of optimal. So take a covering path $\mathcal{P} = (S_1, S_2, \ldots, S_k)$. Let i^* be the machine with the largest makespan at state S_k and let the load this machine be L^*. Consider the last job j^* that was scheduled on machine i, and let the schedule after scheduling j^* be S_t. Ignoring job j^*, at time t the makespan of all the machines is at least $L^* - p_{j^*}$. If not, job j^* would not have been scheduled at machine i^*. Consequently, we have $\sum_{1 \le j \le n} p_i \ge m(L^* - p_{j^*})$. Thus, if OPT is the value of the optimal schedule, then $\text{OPT} \ge \sum_{1 \le j \le n} p_j / m \ge L^* - p_{j^*}$. Clearly $\text{OPT} \ge p_{j^*}$ and so $L^* = L - p_{j^*} + p_{j^*} \le 2\text{OPT}$.

3 Related Work

Here we give a brief overview of related work in this area. The consequences of selfish behavior and the question of efficient computation of Nash equilibria have recently drawn much attention in computer science [8, 7]. Moreover, the use of the price of anarchy [8] as a measure of the cost of the lack of coordination in a game is now widespread, with a notable success in this realm being the selfish routing game [11]. Roughgarden and Tardos [10] also generalize their results on selfish routing games to non-atomic congestion games. A basic result of Rosenthal [9] defines congestion games for which pure strategy Nash equilibria exist. Congestion games belong to the class of potential games [6] for which any best-response path converges to a pure Nash equilibrium. Milchtaich [5] studied player-specific congestion games and the length of best-response paths in this set of games. Even-Dar et. al. [2] considered the convergence time to Nash equilibria in variants of a load balancing game. They bound the number of required steps to reach a pure Nash equilibrium in these games. Recently, Fabrikant et. al. [3] studied the complexity of finding a pure strategy Nash equilibrium in general congestion games. Their PLS-completeness results show that in some congestion games (including network congestion games) the length of a best-response path in the state graph to a pure Nash equilibrium might be exponential. Goemans et. al. [4] considered market sharing games in modeling a decentralized content distribution policy in ad-hoc networks. They show that the market sharing game is a special case of valid-utility games and congestion games. In addition, they give improved bounds for the price of anarchy in some special cases, and present an algorithm to find the pure strategy Nash equilibrium in the uniform market sharing game. The results of Section 5 extend their results.

4 Basic-Utility and Valid-Utility Games

In this section we consider valid-utility games. First we present results concerning the quality of states at the end of one-round paths. Then we give negative results concerning the non-convergence of k-covering paths.

4.1 Convergence

We use the notation from [13]. In particular, a strategy state is denoted by $S = \{s_1, s_2, \ldots, s_k\} \in \mathcal{S}$. Here s_i is the strategy of player i, where $s_i \subseteq V_i$ and V_i is a groundset of elements (with each element corresponding to an action for player i); $\emptyset_i$ corresponds to a null strategy. We also let $S \oplus s'_i = \{s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_k\}$, i.e. the strategy set obtained if agent i changes its strategy from s_i to s'_i. The social value of a state S is $\gamma(S)$, where γ is a submodular function on the groundset $\cup_i V_i$. For simplicity, in this section we will assume that γ is non-decreasing. Similar results, however, do hold in the general case.

We also denote by $\alpha_i(S)$ the private return to player i from the state S, and we let $\gamma'_{s_i}(S) = \gamma(S \cup s_i) - \gamma(S)$. Thus, formally, the second and third conditions in definition 2 are $\alpha_i(S) \geq \gamma'_{s_i}(S \oplus \emptyset_i)$ and $\sum_i \alpha_i(S) \leq \gamma(S)$, respectively. Of particular interest is the subclass of valid-utility games where we always have $\alpha_i(S) = \gamma'_{s_i}(S \oplus \emptyset_i)$; these games are called *basic-utility games* (examples of which include competitive facility location games).

Theorem 1. *In basic-utility games, the social value of a state at the end of a one-round path is at least $\frac{1}{3}$ of the optimal social value.*

Proof. Let $\Omega = \{\sigma_1, \ldots, \sigma_n\}$ denote the optimum state, and let $T = \{t_1, \ldots, t_n\}$ and $S = \{s_1, \ldots, s_n\}$ be the initial state and final states on the one-round path, respectively; we assume the agents play best response strategies in the order $1, 2, \ldots, n$. So $T^i = \{s_1, \ldots, s_{i-1}, t_i, \ldots, t_n\}$ is a state in our one-round path $\mathcal{P} = \{T = T^1, T^2, \ldots, T^{n+1} = S\}$. Thus, by basicness and the fact that the players use best response strategies, we have $\sum_i \alpha_i(T^{i+1}) = \sum_{i=1}^n \gamma'_{s_i}(T^i \oplus \emptyset_i) \geq \sum_i \gamma'_{\sigma_i}(T^i \oplus \emptyset_i)$. It follows by submodularity that $\sum_i \alpha_i(T^{i+1}) \geq \sum_i \gamma'_{\sigma_i}(S \cup T^i \oplus \emptyset_i) \geq \gamma(\Omega) - \gamma(S \cup T) \geq \gamma(\Omega) - \gamma(S) - \gamma(T)$. Moreover, by basicness, $\gamma(S) - \gamma(T) = \sum_{i=1}^n \gamma(T^{i+1}) - \gamma(T^i) = \sum_i \gamma(T^{i+1}) - \gamma(T^i \oplus \emptyset_i) - \sum_i \gamma(T^i) - \gamma(T^i \oplus \emptyset_i) = \sum_i \gamma'_{s_i}(T^i \oplus \emptyset_i) - \sum_i \gamma'_{t_i}(T^i \oplus \emptyset_i) = \sum_i \alpha_i(T^{i+1}) - \sum_i \gamma'_{t_i}(T^i \oplus \emptyset_i)$. Let $\bar{T}^i = \{\emptyset_1, \ldots, \emptyset_{i-1}, t_i, \ldots, t_n\}$. Then, by submodularity, $\gamma(S) - \gamma(T) \geq \sum_i \alpha_i(T^{i+1}) - \sum_i \gamma'_{t_i}(\bar{T}^i \oplus \emptyset_i) = \sum_i \alpha_i(T^{i+1}) - \gamma(T)$. Hence, $\gamma(S) - \gamma(T) \geq \gamma(\Omega) - \gamma(S) - 2\gamma(T)$. Since $\gamma(S) \geq \gamma(T)$, it follows that $3\gamma(S) \geq$ OPT. □

We suspect this result is not tight and that a factor 2 guarantee is possible. Observe, though, that the above proof gives this guarantee for the special case in which the initial strategy state is $T = \emptyset$.

Theorem 2. *In basic-utility games, the social value of a state at the end of a one-round path beginning at $T = \emptyset$ is at least $\frac{1}{2}$ of the optimal social value and this bound is tight.* □

It is known that any Nash equilibria in any valid-utility game has value within a factor 2 of optimal. So here after just one round in a basic-utility game we obtain a solution which matches this guarantee. However for non-basic-utility games, the situation can be different. We can only obtain the following guarantee, which is tight to within a constant factor.

Theorem 3. *In general valid-utility games, the social value of some state on any one-round path is at least $\frac{1}{2n}$ of the optimal social value.*

Proof. Let $\gamma(\Omega) = \text{OPT}$ and assume that $\gamma(t_1, t_2, \ldots, t_n) \leq \frac{1}{2n}\text{OPT}$. Again, agent i changes its strategy from t_i to s_i given the collection of strategies $T^i = \{s_1, \ldots, s_{i-1}, t_i, \ldots, t_n\}$. If at any state in the path $\mathcal{P} = \{T = T^1, T^2, \ldots, T^{n+1} = S\}$ we have $\alpha_i(T^{i+1}) \geq \frac{1}{2n}\text{OPT}$ then we are done. To see this note that $\alpha_j(T^{i+1}) \geq \gamma(T^{i+1}) - \gamma(T^i \oplus \emptyset_j) \geq 0$, since γ is non-decreasing. Thus $\gamma(T^{i+1}) \geq \sum_j \alpha_j(T^{i+1}) \geq \alpha_i(T^{i+1}) \geq \frac{1}{2n}\text{OPT}$. Hence we have, $\gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) - \gamma(T) = \sum_{j=1}^{i} \gamma(s_1 \cup t_1, \ldots, t_j \cup s_j, t_{j+1}, \ldots, t_n) - \gamma(s_1 \cup t_1, \ldots, s_{j-1} \cup t_{j-1}, t_j, t_{j+1}, \ldots, t_n) \leq \sum_{j=1}^{i} \gamma(T^j \cup s_j) - \gamma(T^j) \leq \sum_{j=1}^{i} \gamma(T^{j+1}) - \gamma(T^{j+1} \oplus \emptyset_j) \leq \sum_{j=1}^{i} \alpha_j(T^{j+1}) < \frac{i}{2n}\text{OPT}$. Consequently $\gamma(\sigma_1 \cup t_1 \cup s_1, \ldots, \sigma_i \cup t_i \cup s_i, \sigma_{i+1} \cup t_{i+1}, \ldots, \sigma_n \cup t_n) - \gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \geq \text{OPT} - \gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \geq \text{OPT} - \gamma(S) - \frac{i}{2n}\text{OPT} \geq \frac{2n-i-1}{2n}\text{OPT}$. Thus, there is a $j > i$ such that $\gamma'_{\sigma_j}(T^{i+1}) \geq \gamma'_{\sigma_j}(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \geq \frac{2n-2i-1}{2n(n-i)}\text{OPT} \geq \frac{1}{2n}\text{OPT}$. Therefore we must obtain $\alpha_j(T^{j+1}) \geq \frac{1}{2n}\text{OPT}$ for some $j > i$. □

4.2 Cyclic Equilibria

Here we show that Theorem 3 is essentially tight, and discuss the consequences of this. Specifically, there is the possibility of convergence to low quality states in games in which every Nash equilibria is of high quality.

Theorem 4. *There are valid-utility games in which every solution on a k-covering path has social value at most $\frac{1}{n}$ of the optimal solution.*

Proof. We consider the following n-player game. The groundset of player i consists of three elements x_i, x'_i and y_i. Let $X = \cup_i x_i$ and $X' = \cup_i x'_i$. We construct a non-decreasing, submodular social utility function in the following manner. For each agent $1 \leq i \leq n$, we have

$$\gamma'_{x_i}(S) = \begin{cases} 1 & \text{if } S \cap (X \cup X') = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

We define $\gamma'_{x'_i}(S)$ in an identical manner. Finally, for each agent $1 \leq i \leq n$, we let $\gamma'_{y_i}(S) = 1, \forall S$. Clearly, the social utility function γ is non-decreasing. To see that it is submodular, it suffices to consider any two sets $A \subseteq B$. If $\gamma'_{x_i}(B) = 1$ then $\gamma'_{x_i}(A) = 1$. This follows as $B \cap (X \cup X') = \emptyset$ implies that $A \cap (X \cup X') = \emptyset$. Hence $\gamma'_{x_i}(A) \geq \gamma'_{x_i}(B)$, $\forall i, \forall A \subseteq B$. Similarly $\gamma'_{x'_i}(A) \geq \gamma'_{x'_i}(B)$, $\forall i, \forall A \subseteq B$. Finally, $\gamma'_{y_i}(A) = \gamma'_{y_i}(B) = 1$, $\forall i, \forall A \subseteq B$. It is well known that a function f is submodular if and only if $A \subseteq B$ implies $f'_j(A) \geq f'_j(B)$, $\forall j \in V - B$. Thus γ is submodular.

With this social utility function, we construct a valid utility system. To do this, we create private utility functions α_i using the following rule (except for a few cases given below), where $X_i = S \cap (x_i \cup x'_i)$.

$$\alpha_i(S) = \begin{cases} 1 + \frac{|X_i|}{|(X \cup X') \cap S|} & \text{if } y_i \in S_i \\ \frac{|X_i|}{|(X \cup X') \cap S|} & \text{if } y_i \notin S_i \end{cases}$$

In the following cases, however, we ignore the rule and use the private utilities given in the table.

s_1	s_2	s_3	$\cdots$	s_{n-1}	s_n	$\alpha_1(S)$	$\alpha_2(S)$	$\alpha_3(S)$	$\cdots$	$\alpha_{n-1}(S)$	$\alpha_n(S)$
x_1	x_2	$\emptyset_3$	$\cdots$	$\emptyset_{n-1}$	$\emptyset_n$	0	1	0	$\cdots$	0	0
x_1	x_2	x_3	$\cdots$	$\emptyset_{n-1}$	$\emptyset_n$	0	0	1	$\cdots$	0	0
			$\vdots$						$\vdots$		
x_1	x_2	x_3	$\cdots$	x_{n-1}	$\emptyset_n$	0	0	0	$\cdots$	1	0
x_1	x_2	x_3	$\cdots$	x_{n-1}	x_n	0	0	0	$\cdots$	0	1
x'_1	x_2	x_3	$\cdots$	x_{n-1}	x_n	1	0	0	$\cdots$	0	0
x'_1	x'_2	x_3	$\cdots$	x_{n-1}	x_n	0	1	0	$\cdots$	0	0
			$\vdots$						$\vdots$		
x'_1	x'_2	x'_3	$\cdots$	x'_{n-1}	x_n	0	0	0	$\cdots$	1	0
x'_1	x'_2	x'_3	$\cdots$	x'_{n-1}	x'_n	0	0	0	$\cdots$	0	1
x_1	x'_2	x'_3	$\cdots$	x'_{n-1}	x'_n	1	0	0	$\cdots$	0	0
x_1	x_2	x'_3	$\cdots$	x'_{n-1}	x'_n	0	1	0	$\cdots$	0	0
			$\vdots$						$\vdots$		
x_1	x_2	x_3	$\cdots$	x_{n-1}	x'_n	0	0	0	$\cdots$	1	0

Observe that, by construction, $\sum_i \alpha_i(S) = \gamma(S)$ for all S (including the exceptions). It remains to show that the utility system is valid. It is easy to check that $\alpha_i(S) \geq \gamma(S) - \gamma(S \oplus \emptyset_i) = \gamma'_i(S \oplus \emptyset_i)$ for the exceptions. So consider the "normal" S. If $S_i \cap (x_i \cup x'_i) = \emptyset$, then $\alpha_i(S) = 1$ when $y_i \in S_i$ and $\alpha_i(S) = 0$ otherwise. In both cases $\alpha_i(S) = \gamma'_i(S \oplus \emptyset_i)$. If $S_i \cap (x_i \cup x'_i) \neq \emptyset$ then

$$\gamma'_i(S \oplus \emptyset_i) = \begin{cases} 2 & \text{if } y_i \in S_i \text{ and } (S - S_i) \cap (X \cup X') = \emptyset \\ 1 & \text{if } y_i \in S_i \text{ and } (S - S_i) \cap (X \cup X') \neq \emptyset \\ 1 & \text{if } y_i \notin S_i \text{ and } (S - S_i) \cap (X \cup X') = \emptyset \\ 0 & \text{if } y_i \notin S_i \text{ and } (S - S_i) \cap (X \cup X') \neq \emptyset \end{cases}$$

Consider the first case. We have that $\alpha'_i(S) = 1 + \frac{|X_i|}{|(X \cup X') \cap S|} = 1 + \frac{|X_i|}{|(X \cup X') \cap S_i| + |(X \cup X') \cap (S - S_i)|} = 1 + \frac{|X_i|}{|X_i| + 0} = 2 = \gamma'_i(S \oplus \emptyset_i)$. It is easy to verify that in the other three cases we also have $\alpha_i(S) \geq \gamma'_i(S \oplus \emptyset_i)$. Thus our utility system is valid. It remains to choose which subsets of each players' groundset will correspond to feasible strategies in our game. We simply allow only the singleton elements (and the emptyset) to be feasible strategies. That the set of possible actions for player i are $\mathcal{A}_i = \{\emptyset_i, x_i, x'_i, y_i\}$. Now it is easy to see that an optimal social solution has value n. Any set of strategies of the form $\{y_1, \ldots, y_{i-1}, z_i, y_{i+1}, \ldots, y_n\}$, where $z_i \in \{x_i, x'_i, y_i\}$, $1 \leq i \leq n$ gives a social outcome of value n. However, consider the case of $n = 3$. From the start of the game, if the players behave greedily then we can obtain the sequence of strategies illustrated in Figure 1. The private payoffs given by these exceptional strategy sets mean that each arrow actually denotes a best response move by the labelled agent. However, all of the (non-trivial) strategy sets induce a social outcome of value 1, a factor 3 away from

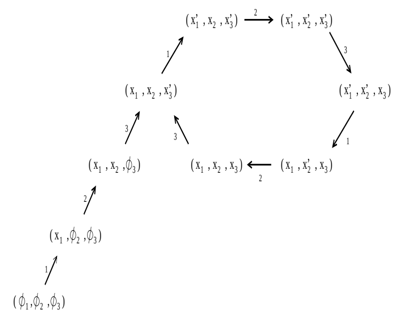

Fig. 1. Bad Cycling.

optimal. Clearly this problem generalizes to n agents. So we converge to a cycle of states all of whose outcomes are a factor n away from optimal. □

So our best response path may lead to a cycle on which *every* solution is extremely bad socially, despite the fact that every Nash equilibria is very good socially (within a factor two of optimal). We call such a cycle in the state graph a *cyclic equilibria*. The presence of low quality cyclic equilibria is therefore disturbing: even if the price of anarchy is low we may get stuck in states of very poor social quality! We remark that our example is unstable in the sense that we may leave the cyclic equilibria if we permute the order in which players make there moves. We will examine in more detail the question of the stability of cyclic equilibria in a follow-up paper.

5 Market Sharing Games

In this section we consider the market sharing game. We are given a set U of n agents and a set H of m markets. The game is modelled by a bipartite graph $G = (H \cup U, E)$ where there is an edge between agent j and market i if market i is of interest to agent j (we write j is interested in market i). The value of a market $i \in H$ is represented by its query rate q_i (this is the rate at which market i is requested per unit time). The cost, to any agent, of servicing market i is C_i. In addition, agent j has a total budget B_j. It follows that a set of markets $s_j \subseteq H$ can be serviced by player j if $\sum_{i \in s_j} C_i \leq B_j$; in this case we say that s_j represents a feasible action for player j. The goal of each agent is to maximise its return from the markets it services. Any agent j receives a reward (return) R_i for providing service to market i, and this reward is dependent upon the number of agents that service this market. More precisely, if the number of agents that serve market i is n_i then the reward $R_i = \frac{q_i}{n_i}$. Observe that the total reward received by all the players is equal to the total query rate of the markets being serviced (by at least one player). The resultant game is called the *market sharing game*. Observe that if $\alpha_j(S)$ is the return to agent j from state S, then the social value is $\gamma(S) = \sum_{j \in U} \alpha_j(S)$. It is then easy to show that the market sharing game is a valid-utility game [4]. We remark that the subcase in which all markets have the same cost is called the *uniform market sharing game*; the subcase in which the bipartite graph G is complete is called the *complete market sharing game*. Note that in this game, the strategy of each player is to solve a knapsack problem. Therefore, in order to model computationally constrained agents, we may assume that the agents apply λ-approximation algorithms to determine their best-response strategies. We then obtain the following theorems concerning the social value after one round of best responses moves.

Theorem 5. *In market sharing games, the social value of a state at the end of a one-round path is at least* $\frac{1}{2H_n+1}$ *of the optimal social value (or at least* $\frac{1}{(\lambda+1)H_n+1}$ *if the agents use* λ*-approximation algorithms).*

Proof. Let $\Omega = \{\sigma_1, \ldots, \sigma_n\}$ denote an optimum state. Here $\sigma_j \subseteq H$ is the set of markets that player j services in this optimum solution; we may also assume that each market is provided by at most one player. Let $T = \{t_1, \ldots, t_n\}$ and $S = \{s_1, \ldots, s_n\}$ be the initial state and final states on the one-round path, respectively. Again, we assume the agents play best response strategies in the order $1, 2, \ldots, n$. So in step r, using a λ-approximation algorithm, agent r changes its strategy from t_r to s_r; thus $T^r = \{s_1, \ldots, s_r, t_{r+1}, \ldots, t_n\}$ is an intermediate state in our one-round path $\mathcal{P} = \{T = T^0, T^1, \ldots, T^n = S\}$. Let $\alpha_j(S)$ be the return to agent j, then the social value of the state $S = T^n$ is $\gamma(S) = \sum_{j \in U} \alpha_j(T^n)$ So we need to show that $\sum_{j \in U} \alpha_j(T^n) \geq \frac{1}{1+(\lambda+1)H_n}\mathrm{OPT}$. Towards this goal, we first show that $\gamma(S) = \sum_{j \in U} \alpha_j(T^n) \geq \frac{1}{H_n} \sum_{j \in U} \alpha_j(T^j)$. We know that agent j does not changes its strategy from s_r after step r. Therefore a market i has a nonzero contribution in $\gamma(S)$ if and only if market i has a nonzero contribution in the summation $\sum_{j \in U} \alpha_j(T^j)$. For any market i, if i appears in one of strategies in T^n then the contribution of i to $\gamma(S)$ is q_i. On the other hand, at most n players use market i in their strategies. Consequently, the contribution of market i in the summation $\sum_{j \in U} \alpha_j(T^j)$ is at most $(1+\frac{1}{2}+\frac{1}{3}+\ldots+\frac{1}{n})q_i = H_n q_i$. It follows that $\sum_{j \in U} \alpha_j(T^n) \geq \frac{1}{H_n} \sum_{j \in U} \alpha_j(T^j)$, as required. We denote by $\mathcal{T}$ the summation $\sum_{j \in U} \alpha_j(T^j)$. Next consider the optimal assignment Ω, and let Y_j be the set of markets that are in serviced by agent j in σ_j but that are *not* serviced by any agent in T^n, that is, $Y_j = \sigma_j - \cup_{r \in U} s_r$. Now $\gamma(S)$ is greater than the value of all the markets in $\cup_{r \in U}(\sigma_r - Y_r)$ since these markets are a subset of markets serviced in T^n. Hence, using the notation $q(Q) = \sum_{i \in Q} q_i$ to denote the sum of query rates of a subset Q of the markets, we have $\gamma(S) \geq \sum_{r \in U} q(\sigma_r - Y_r)$. Next we will prove that $\mathcal{T} \geq \frac{1}{\lambda} \sum_{r \in U} q(Y_r)$. Let Y'_j be the markets in Y_j that are serviced in T^j, that is, $Y'_j = Y_j - (s_1 \cup \cdots \cup s_j \cup t_{j+1} \cup \cdots \cup t_n)$. Then Y'_j is a feasible strategy for agent j at step j, and thus, since player j uses a λ-approximation algorithm, we have $\lambda \alpha_j(T^j) \geq q(Y'_j)$. Therefore, $\lambda \mathcal{T} \geq \sum_{r \in U} q(Y'_r)$.

Finally, we claim that $\mathcal{T} \geq \sum_{j \in U} q(Y''_j)$. To see this, consider a any market $i \in Y''_j = Y_j - Y'_j$. Then market i is not in the strategy set of any agent in T^n, but is in the strategy set of at least one player in T^j. Therefore, somewhere on the path $\mathcal{P}$ after T^j some player must change its strategy and discontinue servicing market i. Let b_i be time step such that T^{b_i} is the first state amongst $T^{j+1}, \ldots, T^n$ that does not service market i. Let $M_j = \{i \in H | b_i = j\}$ be the set of markets for which $b_i = j$. It follows that $\cup_{r \in U} Y''(t) = \cup_{r \in U} M_r$. Notice that $M_r \subseteq t_r$ and no other agents service any market in M_r at step r. It follows that $\alpha_j(T^j) \geq q(M_j)$. Therefore, $\sum_{j \in U} q(Y_j'') = \sum_{j \in U} q(M_j) \leq \sum_{j \in U} \alpha_j(T^j) = \mathcal{T}$. Hence we have, $\text{OPT} = \sum_{j \in U} q(\sigma_j) \leq \sum_{j \in U} q(\sigma_j - Y_j) + \sum_{j \in U} q(Y_j) \leq \gamma(S) + \sum_{j \in U} q(Y_j') + \sum_{j \in U} q(Y_j'') \leq \gamma(S) + \lambda \mathcal{T} + \mathcal{T} \leq (1 + (\lambda + 1) H_n) \gamma(S)$. □

Theorem 6. *In market sharing games, the social value of a state at the end of a one-round path may be as bad as $\frac{1}{H_n}$ of the optimal social value. In particular, this is true even for uniform market sharing games and complete market sharing games.*

Proof. Consider the following instance of a complete market sharing game. There are $m = n$ markets, and the query rate of market i is $q_i = \frac{n}{i} - \epsilon$ for all $1 \leq i \leq n$ where ϵ is sufficiently small. The cost of market i is $C_i = 1 + (n-i)\epsilon$ for $2 \leq i \leq n$ and $C_1 = 1$. There are n players and the budget of player j is equal to $1 + (n-j)\epsilon$. Consider the ordering $1, 2, \ldots, n$ and the one-round path starting from empty set of strategies and letting each player play once in this order. The resulting assignment after this one-round path is that all players provide market number 1 and the social value of this assignment is $n - \epsilon$. However, the optimum solution is for agent j to service market j giving an optimal social value of $nH_n - n\epsilon$. Thus, the ratio between the optimum and the value of the resulting assignment is H_n at the end of a one-round path.

The bad instance for the uniform market sharing game is similar. There are n markets of cost 1 and query rates $q_i = \frac{n}{i} - \epsilon$ for all $1 \leq i \leq n$ where ϵ is sufficiently small. There are n players each with budget 1. Player j is interested in markets $j, j+1, \ldots, n$ and in market 1. It follows that the social value of the assignment after one round is $n - \epsilon$. The optimal social value covers all markets and its value is $nH_n - n\epsilon$. Thus, the ratio is $\frac{1}{H_n}$ after one round. □

6 Conclusion and open problems

In this paper, we presented a framework for studying speed of convergence to approximate solutions in competitive games. We proved bounds on the outcome of one round of best responses of players in terms of the social objective function. More generally, one may consider longer (but polynomial-sized) best-response paths, provided the problem of cycling can be dealt with. In acyclic state graphs, such as potential games (or congestion games), the PLS-completeness results of Fabrikant et. al. [3] show that there are games for which the size of the shortest best-response path from some states to any pure Nash equilibrium is exponential. This implies that in some congestion games the social value of a state after exponentially many best responses might be far from the optimal social value. However, this does not preclude the possibility that good approximate solutions are obtained when short k-covering paths are used. This provides additional motivation for the study of such paths. Here we may consider using a local optimization algorithm and evaluating the output of this algorithm after a polynomial number of local improvements.

The market sharing games are not yet well understood. In particular, it is not known whether exponentially long best-response paths may exist. Bounding the social value of a vertex at the end of a k-covering path is another open question. Goemans et.al. [4] give a polynomial-time algorithm to find the pure Nash equilibrium in uniform market sharing games. Finding such an equilibrium is NP-complete for the general case, but the question of obtaining approximate Nash equilibria is open.

Acknowledgments. This work began at the Bell-Labs Algorithms Pow-Wow; we thank Bruce Shepherd for hosting us at Bell-Labs. The first author would also like to thank Michel Goemans for useful discussions.

References

1. P. Brucker, J. Hurink, and F. Werner. "Improving local search heuristics for some scheduling problems", *Discrete Applied Mathematics*, **72**, pp47-69, 1997.
2. E. Even-dar, A. Kesselman, and Y. Mansour. "Convergence time to Nash equilibria", *ICALP*, 2003.
3. A. Fabrikant, C. Papadimitriou, and K. Talwar. "On the complexity of pure equilibria", *STOC*, 2004.
4. M. Goemans, L. Li, V.S.Mirrokni, and M. Thottan. "Market sharing games applied to content distribution in ad-hoc networks", *MOBIHOC*, 2004.
5. I. Milchtaich. "Congestion games with player-specific payoff functions", *Games and Economic Behavior*, **13(1)**, pp111-124, 1996.
6. D. Monderer and L. Shapley. "Potential games", *Games and Economic Behavior*, **14(1)**, pp124-143, 1996.
7. N. Nisan and A. Ronen. "Algorithmic mechanism design", *Games and Economic Behavior*, **35(1-2)**, pp166–196, 2001.
8. C. Papadimitriou. "Algorithms, games, and the internet", *STOC*, 2001.
9. R. W. Rosenthal. "A class of games possessing pure-strategy Nash equilibria", *International Journal of Game Theory*, **2**, pp65-67, 1973.
10. T. Roughgarden and E. Tardos. "Bounding the inefficiency of equilibria in nonatomic congestion games", *Games and Economic Behavior*, **47(2)**, pp389-403, 2004.
11. T. Roughgarden and E. Tardos. "How bad is selfish routing?", *J. ACM*, **49(2)**, pp236-259, 2002.
12. P. Schuurman and T. Vredeveld. "Performance guarantees of local search for multiprocessor scheduling", *IPCO*, 2001.
13. A. Vetta. "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions", *FOCS*, 2002.