Convergence Issues in Competitive Games

Vahab S. Mirrokni[1] , Adrian Vetta[2] and A. Sidoropoulous [3]

[1]Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT. Email: mirrokni@theory.lcs.mit.edu
[2]School of Computer Science, and Department of Mathematics & Statistics, McGill University. Email: vetta@math.mcgill.ca
[3]Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT. Email: tasos@theory.lcs.mit.edu

Abstract

This report contains the following:

- V. S. Mirrokni, A. Vetta, Convergence Issues in competitive Games, RANDOM-APPROX 2004.
- V. S. Mirrokni, A. Sidoropoulous, Convergence in Cut Games, Submitted.
- The abstract of a working paper by V. S. Mirrokni, and A. Vetta, Sink equilibria and Convergence, In preparation.

Chapter 1

Convergence Issues in Competitive Games

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 [20]. 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.

1.1 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 [13]; 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 1.4.2, 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 2.2, we describe the problem formulations and justify our model. In section 1.3, we present some related work and their relation to this paper. In section 1.4, 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 1.5, 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 1.6, we discuss other classes of games and present some open questions.

1.2 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 $f(\emptyset) = 0$ and, 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 participate.*

- *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 [20]. 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 1.5.

1.2.1 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. [4] 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 [2]. In addition, from any state there is a path of length at most n to some pure Nash equilibrium [19]. 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 $\sum_{1\leq j\leq n} p_i \geq m(L^* - p_{j^*})$. Thus, if OPT is the value of the optimal schedule, then $\mathrm{OPT} \geq \sum_{1\leq j\leq n} p_j/m \geq L^* - p_{j^*}$. Clearly $\mathrm{OPT} \geq p_{j^*}$ and so $L^* = L - p_{j^*} + p_{j^*} \leq 2\mathrm{OPT}$.

1.3 Related Work

In this section 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 [13, 12]. Moreover, the use of the price of anarchy [13] 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 [17]. Roughgarden and Tardos [15] also generalize their results on selfish routing games to non-atomic congestion games.

A basic result of Rosenthal [16] defines congestion games for which pure strategy Nash equilibria exist. In these games, each player chooses a particular combination of factors from a common set; the payoff from each factor is then a function of the number of players who chose this factor. Congestion games belong to the class of potential games [11] for which any best-response path converges to a pure Nash equilibrium. Milchtaich [9] studied player-specific congestion games and the length of best-response paths in this set of games.

Even-Dar et. al. [4] 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. They consider a central policy that let agents move in a certain order, and analyze different policies to choose such an ordering. In particular, they consider the random and FIFO model. In contrast to this work, their interest is in the convergence time to a pure Nash equilibrium and not to good approximate solutions.

Recently, Fabrikant et. al. [5] 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. [6] 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 1.5 extend their results.

1.4 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.

1.4.1 Convergence

We use the notation from [20]. 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

$\sum_i \alpha_i(S) \le \gamma(S)$, 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*.

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. 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 $1, 2, \ldots, n$. Thus $T^i = \{s_1, \ldots, s_{i-1}, t_i, \ldots, t_n\}$ is an intermediate 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

$$\begin{aligned} \sum_i \alpha_i(T^{i+1}) &= \sum_{i=1}^n \gamma'_{s_i}(T^i \oplus \emptyset_i) \\ &\ge \sum_i \gamma'_{\sigma_i}(T^i \oplus \emptyset_i) \end{aligned}$$

It follows by submodularity that

$$\begin{aligned} \sum_i \alpha_i(T^{i+1}) &\ge \sum_i \gamma'_{\sigma_i}(S \cup T^i \oplus \emptyset_i) \\ &\ge \gamma(\Omega) - \gamma(S \cup T) \\ &\ge \gamma(\Omega) - \gamma(S) - \gamma(T) \end{aligned}$$

Moreover, by basicness,

$$\begin{aligned} \gamma(S) - \gamma(T) &= \sum_{i=1}^n \gamma(T^{i+1}) - \gamma(T^i) \\ &= \sum_i \gamma(T^{i+1}) - \gamma(T^i \oplus \emptyset_i) - \sum_i \gamma(T^i) - \gamma(T^i \oplus \emptyset_i) \\ &= \sum_i \gamma'_{s_i}(T^i \oplus \emptyset_i) - \sum_i \gamma'_{t_i}(T^i \oplus \emptyset_i) \\ &= \sum_i \alpha_i(T^{i+1}) - \sum_i \gamma'_{t_i}(T^i \oplus \emptyset_i) \end{aligned}$$

Let $\bar{T}^i = \{\emptyset_1, \ldots, \emptyset_{i-1}, t_i, \ldots, t_n\}$. Then, by submodularity,

$$\begin{aligned} \gamma(S) - \gamma(T) &\ge \sum_i \alpha_i(T^{i+1}) - \sum_i \gamma'_{t_i}(\bar{T}^i \oplus \emptyset_i) \\ &= \sum_i \alpha_i(T^{i+1}) - \gamma(T) \end{aligned}$$

Hence,

$$\gamma(S) - \gamma(T) \ge \gamma(\Omega) - \gamma(S) - 2\gamma(T)$$

Since $\gamma(S) \ge \gamma(T)$, it follows that $3\gamma(S) \ge \text{OPT}$. □

We suspect this result is not tight and that a factor 2 guarantee is possible. This guarantee is obtained 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.*

Proof. Using $\gamma(T) = 0$, the result follows from the last inequlity of the proof of Theorem 2. $\square$

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 performance 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

$$\begin{aligned}
\gamma(T^{i+1}) &\geq \sum_j \alpha_j(T^{i+1}) \\
&\geq \alpha_i(T^{i+1}) \\
&\geq \frac{1}{2n}\text{OPT}
\end{aligned}$$

Hence we have,

$$\begin{aligned}
&\gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) - \gamma(T) \\
&\quad= \sum_{j=1}^{i} \gamma(s_1 \cup t_1, \ldots, t_j \cup s_j, t_{j+1}, \ldots, t_n) - \gamma(s_1 \cup t_1, \ldots, s_{j-1} \cup t_{j-1}, t_j, t_{j+1}, \ldots, t_n) \\
&\quad\leq \sum_{j=1}^{i} \gamma(T^j \cup s_j) - \gamma(T^j) \\
&\quad\leq \sum_{j=1}^{i} \gamma(T^{j+1}) - \gamma(T^{j+1} \oplus \emptyset_j) \\
&\quad\leq \sum_{j=1}^{i} \alpha_j(T^{j+1}) \\
&\quad< \frac{i}{2n}\text{OPT}
\end{aligned}$$

It follows that

$$\begin{aligned}
&\gamma(\sigma_1 \cup t_1 \cup s_1, \ldots, \sigma_i \cup t_i \cup s_i, \sigma_{i+1} \cup t_{i+1}, \ldots, \sigma_n \cup t_n) - \gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \\
\geq\ & \text{OPT} - \gamma(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \\
\geq\ & \text{OPT} - \gamma(S) - \frac{i}{2n}\text{OPT} \\
\geq\ & \frac{2n - i - 1}{2n}\text{OPT}
\end{aligned}$$

Thus, there is a $j > i$ such that

$$\gamma'_{\sigma_j}(T^{i+1}) \geq \gamma'_{\sigma_j}(t_1 \cup s_1, \ldots, t_i \cup s_i, t_{i+1}, \ldots, t_n) \geq \frac{2n - 2i - 1}{2n(n-i)}\text{OPT} \geq \frac{1}{2n}\text{OPT}$$

Consequently, we must obtain $\alpha_j(T^{j+1}) \geq \frac{1}{2n}\text{OPT}$ for some $j > i$. $\square$

1.4.2 Cyclic Equilibria

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

Theorem 4 *There are valid-utility games that are not basic-utility in which every solution on a k-covering path has social value at most $\frac{1}{n}$ of the optimal solution, for all $k > 0$.*

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

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

Similarly, we have

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

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
x'_1	x'_2	x'_3	$\cdots$	x_{n-1}	x_n	0	0	1	$\cdots$	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
x_1	x_2	x_3	$\cdots$	x'_{n-1}	x'_n	0	0	1	$\cdots$	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 is for player i, we have $\mathcal{A}_i = \{\emptyset_i, x_i, x_i', y_i\}$. Now it is easy to see that an optimal social solution has value n. Any collection of strategies of the form $\{y_1, y_2, \ldots, y_{i-1}, z_i, y_{i+1}, \ldots, y_n\}$, 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 $n = 3$. 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 1.1. The private payoffs given by these exceptional strategy sets mean that each arrow

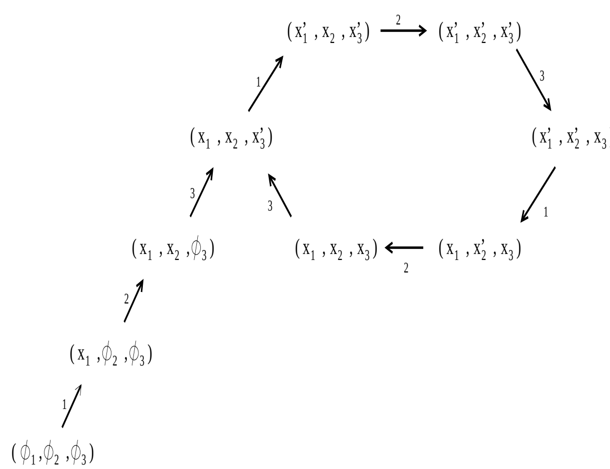

Figure 1.1: Bad Cycling.

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 optimal. Clearly this problem generalizes to n agents. Hence, we converge to a cycle of states all of whose outcomes are a factor n away from optimal. $\square$

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 very 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.

1.5 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 modeled 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 maximize 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 $\alpha_j(S)$ be 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 [6]. 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 $\sum_{j \in U} \alpha_j(T^n) \geq \frac{1}{1+(\lambda+1)H_n}\text{OPT}$. 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,

$$\begin{aligned}
\text{OPT} &= \sum_{j\in U} q(\sigma_j) \\
&\leq \sum_{j\in U} q(\sigma_j - Y_j) + \sum_{j\in U} q(Y_j) \\
&\leq \gamma(S) + \sum_{j\in U} q(Y'_j) + \sum_{j\in U} q(Y''_j) \\
&\leq \gamma(S) + \lambda\mathcal{T} + \mathcal{T} \\
&\leq (1 + (\lambda+1)H_n)\gamma(S)
\end{aligned}$$

□

In the following theorem we prove that the above result is tight upto a constant factor.

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. □

In the above tight examples, the strategy of the agents is not symmetric. Here, we observe that in the uniform market sharing game if the strategy of all players is *symmetric*, i.e., the cost of all markets is the same and each player can provide a fixed, say k, number of markets then the social value of a vertex at the end of a covering path is at least $\frac{1}{2}$ of the optimal social value. The reason is that in this symmetric variant of the market sharing game, any vertex at the end of a covering path is a pure Nash equilibrium. In fact after a player j plays his best response, he will not have incentive to change his strategy in the game. To prove this claim, we assume for contradiction the first place at which player j wants to switch from a market i to a new market i' and assume that this occurs after player j' moves. A careful case analysis shows that this cannot happen, since either player j' would have had incentive to change to i' instead of i or player j had incentive to switch to i' before player j' moves. The details of this proof is omitted here. Thus we know that

Theorem 7 *In the symmetric uniform market sharing game, the vertex at the end of any covering path is a pure strategy Nash equilibrium. In particular, the social value of this vertex is at least $\frac{1}{2}$ of the optimal social value.*

1.6 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. [5] 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. Consider, for example, example the maximum weighted cut problem[1]. There are examples in which the value of the Max-Cut is $O(\frac{1}{n})$ the value of the maximum cut after a constant number c of rounds, that is at the end of a c-covering path. Consequently, one problem is to bound the cost of the resultant cut after 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. [6] 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 would like to thank Bruce Shepherd for hosting us at Bell-Labs. The first author would also like to thank Michel Goemans for useful discussions.

[1]The local optimization algorithm for the Max-Cut problem can be formalized as a congestion game with the size of the cut as the social function. It is known that it is PLS-complete to find the local optimum of the Max-Cut or the pure equilibria of the corresponding game [18].

Chapter 2

Convergence in Cut Games

absract. We study the speed of convergence to approximate solutions in cut games which are related to party affiliation games [5]. In a cut game, each player corresponds to a node of a graph. Each node's payoff is its contribution in the cut. We consider two social functions: the *value of the cut*, and the *total happiness*, i.e., the weight of the edges in the cut, minus the weight of the remaining edges. For these social functions, we analyze the speed of convergence to approximate solutions. The time complexity of a game is measured in *rounds*, where a round consists of a sequence of movements, with each node playing at least once in each round. In the case where each round consists of the same sequence of movements, we prove $\Omega(n)$ and $\Omega(\sqrt{n})$ lower bounds, for weighted and unweighted graphs respectively, for the number of rounds required in the worst case to converge to a constant-factor cut. For the case where each round might consist of a different sequence of movements, we prove an *exponential* lower bound. We also show that a random ordering of players will converge to a constant-factor cut in expectation in one round. We prove that *mildly greedy* players will converge to a constant factor cut in one round, under any ordering. For the total happiness social function, we show that for unweighted graphs of large girth, if we start from a random configuration, then greedy behavior of players in a random order converges to an approximate solution after one round.

2.1 Introduction

Algorithmic design for systems with the presence of selfishness, has to address the problem of lack of coordination. The major tool for analyzing such systems is the notion of the *price of anarchy* in a game [13]; this is the worst case ratio between an optimal social solution and a Nash equilibrium. A high price of anarchy certainly indicates that the functioning system is in need of central regulation. On the other hand, a low price of anarchy does not necessarily imply high performance of the system. One main reason is that in many games, the repeated selfish behavior of players may not lead to a Nash equilibrium. Moreover, even if the selfish behavior of players converges to a Nash equilibrium, the *speed* of convergence might be very slow. Thus, from a practical and computational viewpoint, it is important to evaluate the speed or rate of convergence to approximate solutions.

If we model the repeated selfish behavior of players as a sequence of improvements by players, the convergence question is related to the running time of local optimization algorithms. In fact, the theory of PLS-completeness [18] and the known exponentially long paths in many local optimization problems such as Max-2SAT and Max-Cut, indicates that in many of these settings, we cannot hope for a polynomial-time convergence to a Nash equilibrium. For these games, it is not sufficient to just study the value of the social function at Nash equilibria. To deal with this issue, we need to bound the social value of a strategy profile after polynomially many improvements by players. We consider (best response) paths in this graph and evaluate the social function at states at the end of these paths. The rate of convergence to high quality solutions (or Nash equilibria) can then be measured by the length of the path.

Related Work The Max-Cut problem is a well-studied problem [7], and local optimization algorithms have been considered for this problem. It is well known that finding a local optimum for Max-Cut, is PLS-complete [8, 18] and there are some configurations from which paths to a local optimum are exponentially long. In the positive side, Poljak [14] proved that for cubic graphs the convergence time is at most $O(n^2)$ steps. The maximum total happiness social function is considered in the context of correlation clustering [1], and is similar to the total agreement minus disagreement in that context. The best approximation algorithm known for this problem is an $O(\log(n))$-approximation algorithm [3] based on a semidefinite programming relaxation.

This work is motivated by the negative results of the convergence for congestion games [5], and the study of convergence in valid-utility and basic-utility games [10]. Fabrikant, Papadimitriou, and Talwar [5] show that for general congestion and asymmetric network congestion games, the problem of finding a pure Nash equilibrium is PLS-complete. This implies exponentially long paths to equilibria for these games. Mirrokni and Vetta [10] address the convergence to approximate solutions on the same model on basic-utility and valid-utility games. They prove that one round of selfish behavior of players, converges to a $\frac{1}{3}$-approximate solution in basic-utility games. For valid-utility games, they show existence of poor cycles in the state graph. The cut game that we consider is a congestion game [16]. It is known that in these games, best-response paths converge to a Nash equilibrium, and pure Nash equilibrium exists. Other related papers that considered convergence for different classes of games are on load balancing games [4], market sharing games [6] and congestion games [5].

Our Contribution. In this paper, we focus on a cut game, motivated from a *party affiliation game* [5], and the local optimization algorithm for the maximum cut problem. In the party affiliation game, each player's strategy is to choose one of two parties, i.e, $s_i \in \{1, -1\}$ and the payoff of player i for the strategy profile $(s_1, s_2, \ldots, s_n)$ is $\sum_j s_j s_i w_{ij}$. It is well-known that the pure Nash equilibrium exists in this game, and selfish behavior of players converges to a Nash equilibrium, since this game is a potential game. We can model this game as the following cut game: each node of a graph is a player and a node's payoff is his contribution in the cut. Thus a player moves if he can improve his contribution in the cut or equivalently he can improve the value of the cut. We consider two social functions: the value of the cut and the value of the cut minus the value of the rest of edges. Our work deviates from bounding the distance to a Nash equilibrium [18, 5], and focus in studying the rate of convergence to an approximate solution [10] for cut games. The time complexity of a game is measured in *rounds*, where a round consists of a sequence of movements, with each node playing at least once in each round.

On the negative side, we prove the existence of paths with *exponential* number of rounds, that are poor in terms of the value of the resulting cut. In the case where each round consists of the same sequence of movements, we prove $\Omega(n)$ and $\Omega(\sqrt{n})$ lower bounds, for weighted and unweighted graphs respectively, for the number of rounds required in the worst case to converge to a constant-factor cut. On the positive side, we prove that selfish behavior of players in a round in which the ordering of the player is picked uniformly at *random*, the value of the cut is a $\frac{1}{8}$-approximation in expectation. We also model the selfish behavior of *mildly greedy* players that move if their payoff increases by a factor $1+\epsilon$. We prove that in contrast to totally greedy players, mildly greedy players converge to a constant-factor cut after one round. For unweighted graphs, we prove that they converge to a constant-factor cut after at most $O(n)$ rounds. Finally, for the total happiness social function, we show that for unweighted graphs of large girth, if we start from a random configuration, then greedy behavior of players in a random order converges to an approximate solution after one round.

The rest of the paper is organized as follows. In Section 2.2, we mention all definitions and preliminaries. In Section 2.3, we prove our negative results for the convergence for the cut value. Our positive results for fast convergence to constant-factor cuts are described in Section 2.4. In Section 2.5, we study the convergence to approximate solutions for maximum happiness, after a random configuration and a random ordering of players. In Section 2.6, we conclude the paper by discussing the usefulness of the model in game theory and local optimization and with some open questions.

2.2 Definitions and Preliminaries

We consider the following cut game. We are given an undirected graph $G(V, E)$ with n vertices and edge weights $w : E(G) \rightarrow \mathcal{R}$. Each node of the graph is a player. Each player's action or strategy is to choose one side of the cut, i.e, node v can choose $s_v = -1$ or $s_v = 1$. An action profile corresponds to a cut in the graph. The *payoff* of player v is equal to his total contribution in the cut, i.e, the total weight of the edges in the neighborhood of v in the cut ($p(v) = \sum_{j:s_j \neq s_v} w_{jv}$). The *happiness* of a node v is equal to his contribution in the cut minus the value of the edges in the same side as this node. We consider two social functions: the cut value and the cut value minus the value of the rest of the edges in the graph. It is easy to see that the cut value is half the sum of payoffs of nodes. The second social function is half the sum of happiness of nodes. We call the second social function, *total happiness*.

In order to model the selfish behavior of players, we use the underlying *state graph*. Each vertex in the state graph represents a strategy state $S = (s_1, s_2, \ldots, s_n)$ and corresponds to a cut. The arcs in the state graph corresponds to best response moves by the players.

Formally, we have

Definition 4 *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 labeled 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 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 labeled $i_1, i_2, \ldots, i_n$ 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 $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_k)$.

Random one-round path: Let σ be a ordering of players picked uniformly at random from all possible orderings. The one-round path $\mathcal{P}$ corresponding to the ordering σ is a random one-round path.

Random path: A path $\mathcal{P}$ in the state graph is a random path, if at each step we choose the next player who moves uniformly at random.

This model is the same as the model used by Mirrokni and Vetta [10]. 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 the final state after one-round, covering, and k-covering, random, and random one-round paths.

Let $c(G)$ be the value of the maximum cut in graph G. Let $\sigma : \{1, \ldots, |V(G)|\} \rightarrow V(G)$, be an ordering of the vertices of G, and $k > 0$, be an integer. We denote by $\alpha(G, \sigma, k)$, the size of the cut after k of one-round paths in the state graph, with respect to the ordering σ. We also denote by $\alpha^R(G,k)$, the expected value of the cut at the end of a random one-round path. Let also $\beta(G,\sigma,k) = \frac{c(G)}{\alpha(G,\sigma,k)}$, and $\beta^R(G,k) = \frac{c(G)}{\alpha(G,k)}$.

2.3 Lower Bounds for Convergence to a Constant-Factor Cut

In this section, we give lower bounds for the convergence to approximate solutions for the cut social function.

Theorem 8 *There exists a weighted graph $G = (V(G), E(G))$, with $|V(G)| = \Theta(n)$, and a k-covering path $\mathcal{P}$ in the state graph, for some k exponentially large in n, such that the value of the cut at the end of $\mathcal{P}$, is at most $O(1/n)$ of the optimum cut.*

Proof. In [18], it is shown that there exists a weighted graph $G_0 = (V(G_0), E(G_0))$, and an initial cut $(S_0, \bar{S}_0)$, such that the length of *any* path in the state graph, from $(S_0, \bar{S}_0)$ to a Nash equilibrium, is exponentially long. Consider such a graph of size $\Theta(n)$, with $V(G_0) = \{v_0, v_1, \ldots, v_N\}$, and let $\mathcal{P}_0$, be an exponentially long path from $(S_0, \bar{S}_0)$ to a Nash equilibrium. Let $s_0, s_1, \ldots, s_{|\mathcal{P}|}$, be the sequence of states visited by $\mathcal{P}$. The construction of [18] guarantees that for any state s_i, there is exactly one node in G_0, which wants to change his side of the cut. That is, there is a single move which improves the size of the cut. For any $i \in \{1, \ldots, |\mathcal{P}|\}$, let y_i be the node in G_0, which changes side in the cut, while we move from state s_{i-1}, to state s_i.

Since there are $\Theta(n)$ nodes in G_0, we can assume w.l.o.g., that the node $v_0 \in V(G_0)$, changes side in the cut an exponential number of times, along path $\mathcal{P}_0$. Let $t_0 = 0$, and for $i \geq 1$, let t_i be the time in which v_0 changes side for the ith time along the path $\mathcal{P}_0$. For $i \geq 1$, let $\mathcal{Q}_i$ be the sequence of nodes $y_{t_{i-1}+1}, \ldots, y_{t_i}$. Observe that each $\mathcal{Q}_i$, might not contain all of the nodes in G_0. In this case, if $\mathcal{Q}_i = z_1, z_2, \ldots, z_k, v_0$, we can replace it with the sequence $\mathcal{Q}'_i = z_1, z_2, \ldots, z_k, v_1, v_2, \ldots, v_N, v_0$. Clearly, the new sequence contains all the nodes of G_0. Moreover, if we start at state $s_{t_{i-1}}$, and we let the nodes of G_0 play according to the order $\mathcal{Q}'_i$, we will traverse exactly the same path in the state graph, as if we let them play according to $\mathcal{Q}_i$. Thus, in what follows, we may assume w.l.o.g., that each sequence $\mathcal{Q}_i$ contains all the nodes of G_0.

Consider now a graph G, which consists of a path $L = x_1, x_2, \ldots, x_n$, and a copy of G_0. For each $i \in \{1, \ldots, n-1\}$, the weight of the edge $\{x_i, x_{i+1}\}$ is 1. We scale the weights of G_0, such that the total weight of the edges of G_0 is less than 1. Finally, for each $i \in \{1, \ldots, n\}$, we add the edge $\{x_i, v_0\}$, of weight ϵ, for some sufficiently small ϵ. Intuitively, we can pick the value of ϵ, such that the moves made by the nodes in G_0, are independent of the positions of the nodes of the path L in the current cut.

For each $i \geq 1$, we consider an ordering $\mathcal{R}_i$ of the nodes of L, as follows: If i is odd, then $\mathcal{R}_i = x_1, x_2, \ldots, x_n$, and if i is even, then $\mathcal{R}_i = x_n, x_{n-1}, \ldots, x_1$.

We are now ready to describe the exponentially long path in G. Assume w.l.o.g., that in the initial cut for G_0, we have $v_0 \in S_0$. The initial cut for G is $(S, \bar{S})$, with $S = \{x_1\} \cup S_0$, and $\bar{S} = \{x_2, \ldots, x_n\} \cup \bar{S}_0$. It is now straight-forward to verify that there exists an exponentially large k, such that for any i, with $1 \leq i \leq k$, if we let the nodes of G play according to the sequence $\mathcal{Q}_1, \mathcal{R}_1, \mathcal{Q}_2, \mathcal{R}_2, \ldots, \mathcal{Q}_i, \mathcal{R}_i$, then we have:

- If i is odd, then $\{x_1, \ldots, x_{n-1}\} \subset \bar{S}$, and $\{v_0, x_n\} \subset S$.
- If i is even, then $\{v_0, x_1\} \subset S$, and $\{x_2, \ldots, x_n\} \subset \bar{S}$.

It follows that for each i, with $1 \leq i \leq k$, the size of the cut is at most $O(1/n)$ times the value of the optimum cut. The Lemma follows since each path in the state graph induced by the sequence $\mathcal{Q}_i, \mathcal{R}_i$, is a covering path. $\square$

Observe that in the proof of Theorem 8, it is crucial that each covering path corresponds to a different sequence of moves. For the case where each covering path corresponds to the same sequence, we can prove a weaker lower bound. We observed that our example in the following theorem is very similar to an example that Poljak [14] used. He showed that there exist graphs of small degree, that might require $\Omega(n^2)$ moves to converge to a Nash equilibrium, which also implies an $\Omega(n)$ bound for the number of rounds of the same ordering.

Theorem 9 *There exists a weighted graph $G = (V(G), E(G))$, with $|V(G)| = n$, and an ordering of vertices such that for any $k > 0$, the value of the cut after k rounds of letting players play in this ordering is at most $O(k/n)$ of the maximum cut.*

Proof. Consider a graph $G = (V(G), E(G))$, with $V(G) = \{1, 2, \ldots, n\}$, and $E(G) = \bigcup_{i=1}^{i-1}\{\{i, i+1\}\}$. For any i, with $1 \leq i < n$, the weight of the edge $\{i, i+1\}$, is $1 + (i-1)/n^2$. Since G is bipartite, the value of the maximum cut of G is $c(G) = \sum_{i=1}^{n-1}(1 + (i-1)/n^2) = \Omega(n)$.

Let σ, by an ordering of the vertices of G, with $\sigma(i) = i$. Consider the execution of the one-round path for the ordering σ. Initially, we have $S = V(G)$. It is easy to see that in any round $i \geq 1$, when node j plays, if $j \leq n - i$, j moves to the other part of the cut. Otherwise, if $j > n - i$, j remains in the same part of the cut. Thus, after round i, we have

$$S = \begin{cases} \{n, n-2, n-4, \ldots, n-i+1\} & \text{if i is odd} \\ \{1, 2, \ldots, n-i-1)\} \cup \{n, n-2, n-4, \ldots, n-i\} & \text{if i is even} \end{cases}$$

It easily follows, that $\alpha(G, \sigma, k) = \sum_{i=n-k}^{n-1} 1 + (i-1)/n^2 = O(k)$. □

We can apply the same idea to prove a weaker lower bound for unweighted graphs.

Theorem 10 *There exists an unweighted graph $G = (V(G), E(G))$, with $|V(G)| = n$, and an ordering of vertices such that for any $k > 0$, the value of the cut after k rounds of letting players play in this ordering is at most $O(k/\sqrt{n})$ of the maximum cut.*

Proof. Consider the graph a $G = (V(G), E(G))$, with $V(G) = \bigcup_{i=1}^{t}\bigcup_{j=1}^{i}\{\{v_{i,j}\}\}$, and $E(G) = \bigcup_{i=1}^{t-1}\bigcup_{j=1}^{i}\bigcup_{l=1}^{i+1}\{\{v_{i,j}, v$
Clearly, G is bipartite, and thus $c(G) = |E(G)| = \Omega(t^3) = \Omega(n^{3/2})$.

Consider now the ordering σ, such that for any i, j, with $1 \leq j \leq i \leq t$, $\sigma(i(i-1)/2 + j) = v_{i,j}$. By an argument similar to the one used in the proof of Theorem 9, we obtain that $\alpha(G, \sigma, k) = O(kt^2) = O(kn)$. □

2.4 Upper Bounds for Convergence to a Constant-Factor Cut

Now we prove positive results for the convergence to constant-factor approximate solutions. First, we show that the expected value of the cut after a random one-round path is within a constant factor of the maximum cut, i.e, $\beta^R(G, 1) = O(1)$.

Theorem 11 *In weighted graphs, the expected value of the cut at the end of a random one-round path is at least $\frac{1}{8}$ of the maximum cut.*

Proof. We prove that the expected contribution of a node in the cut is greater than $\frac{1}{8}$ of the weight of its adjacent edges. Let $W(v) = \sum_{u \in N(v)} w(u, v)$. The probability that a node v occurs after exactly k of its neighbors is $\frac{1}{\deg(v)}$. Conditioning on the fact that v occurs after k of its neighbors, the probability that a node u, adjacent to v, occurs after v in the ordering is $\frac{\deg(v)-k}{\deg(v)}$. After v moves the contribution of v in the cut is at least $\frac{W(v)}{2}$. Given that v occurs after exactly k neighbors, for each vertex u in the neighborhood of v, the probability that it occurs after v is $\frac{\deg(v)-k}{\deg(v)}$ and only in this case u can decrease the contribution of v in the cut by the value at most $w(u, v)$. Thus the contribution of v in the cut is at least $\max(0, W(v)(\frac{1}{2} - \frac{\deg(v)-k}{\deg(v)}))$. Summing over all values of k, we get that the expected contribution of v in the cut is at least $\sum_{k=0}^{d(v)} \frac{1}{\deg(v)} \max(0, W(v)(\frac{1}{2} - \frac{\deg(v)-k}{\deg(v)})) \geq \frac{W(v)}{8}$. The result follows by the linearity of expectation. □

The above theorem shows that even though for some orderings it takes $\Omega(n)$ rounds to converge to a constant-factor solution, a random ordering gives a constant-factor cut in average. In fact, we can also show that if we pick each vertex at random then after $O(n^2)$ moves we get a constant-factor cut in expectation.

Theorem 12 *The expected value of the cut at the end of a random path of length n^2 is at least $\frac{1}{4}$ of the optimal cut.*

Proof. Let the current cut be $(S, \bar{S})$. Let the current value of the cut be $c'(G)$. Assume that $c'(G)$ is smaller than $\frac{1}{4}$ of the total weight of the edges, W. Otherwise, the cut is already $\frac{1}{4}$ factor of the optimum. Let $\alpha(v)$ be the payoff of vertex v if we let v move at the current cut. We know that $\sum_{v \in V(G)} \alpha(v) \geq W$. We also know that $2c'(G) = \sum_{v \in V(G)} p(v) \leq \frac{W}{2}$. Thus there exists a vertex u for which $\alpha(u) - p(u) \geq \frac{W}{2n}$. This means that at each step there exists a vertex whose movement improves the cut by at least $\frac{W}{2n}$. At each step of the random path, we choose this vertex with probability at least $\frac{1}{n}$. Thus if the current cut is not a $\frac{1}{4}$-approximation, the expected improvement after each step is at least $\frac{W}{4n^2}$. Note that we will either get to a cut with value $\frac{W}{4}$ or the improvement at n^2 steps is at least $\frac{W}{4}$. $\square$

Note that the proof of the above theorem also shows that from any cut there exists a path of length at most n to a $\frac{1}{4}$-approximate cut. Theorem 9 shows that there are graphs for which the value of the cut after any $o(n)$ one-round paths is not a constant-factor cut. Here we observe that if we change the game by assuming that a node change the side if his payoff is multiplied by a factor $1+\epsilon$ for a constant ϵ, then the convergence is faster. We say a node is $1+\epsilon$-greedy if he moves to the other side only if his payoff is multiplied by a factor of $1+\epsilon$. In the following, we prove that if all nodes all $1+\epsilon$-greedy for a constant ϵ, then the cut value after any one-round path is within a constant factor of the optimum.

Theorem 13 *If all nodes are $(1+\epsilon)$-greedy, then the cut value at the end of any one-round path is within a $\min\{\frac{1}{4+2\epsilon}, \frac{\epsilon}{4+2\epsilon}\}$ factor of the optimal cut.*

Proof. Consider a covering path $\mathcal{P}$. For each vertex v, let $\alpha(v)$ be the payoff of v after its occurrence in $\mathcal{P}$. Let $p(v)$ be the payoff of v at the end of $\mathcal{P}$. Let V_1 be the set of vertices that did not change their side in the one-round path and $V_2 = V(G) - V_1$. For a vertex $v \in V_2$, let $u(v)$ be the total weight of the edges that are removed from the cut after v moves. Let $w(v)$ be the total weight of the edges adjacent to vertex v, and for a set $S \subseteq V(G)$, let $w(S) = \sum_{v \in S} w(v)$.

$$\begin{aligned}
\sum_{v \in V(G)} p(v) &= \sum_{v \in V_1} p(v) + \sum_{v \in V_2} p(v) \\
&\geq \sum_{v \in V_1} \alpha(v) + \sum_{v \in V_2} \alpha(v) - \sum_{v \in V_2} u(v) \\
&\geq \frac{1}{2+\epsilon} w(V_1) + \frac{1+\epsilon}{2+\epsilon} w(V_2) - \frac{1}{2+\epsilon} w(V_2) \\
&\geq \min\{\frac{1}{2+\epsilon}, \frac{\epsilon}{2+\epsilon}\} w(V(G))
\end{aligned}$$

Thus the value of the cut is at least $\min(\frac{1}{4+2\epsilon}, \frac{\epsilon}{4+2\epsilon})$-approximation after this round. $\square$

In particular, the above theorem shows that 2-greedy nodes converge to a constant-factor solution after one round and faster than 1-greedy players in the worst case. In other words, less greedy players are socially better in short time in the worst case.

In unweighted simple graphs, it is obvious that the value of the cut at the end of n^2-covering paths is within a constant factor of the optimal cut. The following theorem shows that in unweighted graphs, the value of the cut after any $O(n)$-covering path is a constant-factor approximation.

Theorem 14 *For unweighted graphs, the value of the cut after an n-covering path is within a constant-factor of the maximum cut.*

Proof. Consider a k-covering path $\mathcal{P} = (\mathcal{P}_1, \ldots, \mathcal{P}_k)$. Let $M_0 = 0$, and for any $i \geq 1$, let M_i be the size of the cut at the end of $\mathcal{P}_i$. Note that if $M_i - M_{i-1} \geq \frac{|E(G)|}{10n}$, for all i, with $1 \leq i \leq k$, then clearly $M_k \geq k\frac{|E(G)|}{10n}$, and since $c(G) \leq |E(G)|$, the Lemma follows.

It remains to consider the case where there exists i, with $1 \leq i \leq k$, such that $M_i - M_{i-1} < \frac{|E(G)|}{10n}$. Let V_1 be the set of nodes that change their side in the cut on the path $\mathcal{P}_i$, and let also $V_2 = V(G) \setminus V_1$. Observe that when a node changes its side in the cut, the size of the cut increases by at least 1. Thus, $|V_1| < \frac{|E(G)|}{10n}$, and the number of edges that are incident to at least one node in V_1, is less than $\frac{|E(G)|}{10}$.

On the other hand, if a node of degree d remains in the same part of the cut, then exactly after it plays, at least $\lceil d/2 \rceil$ of its adjacent edges are in the cut. Thus, at least half of the edges that are incident to at least one node in V_2, were in the cut, at some point during path $\mathcal{P}_i$. At most $\frac{|E(G)|}{10}$ of these edges have an end-point in V_1, and thus at most that many of these edges may not appear in the cut at the end of $\mathcal{P}_i$. Thus, the total number of edges that remain in the cut at the end of path $\mathcal{P}_i$, is at least $\frac{|E(G)|-|E(G)|/10}{2} - \frac{|E(G)|}{10} = \frac{7|E(G)|}{20}$. Since $c(G) \leq |E(G)|$, we obtain that at the end of $\mathcal{P}_i$, the value of the cut is within a constant factor of the optimum. $\square$

2.5 The Total Happiness

In this section we show upper bounds for the maximum happiness at the end of a random one-round path, for unweighted graphs of large girth.

First we observe that the price of anarchy is unbounded for this social function. Consider an unweighted cycle of size four, $V(G) = \{v_1, v_2, v_3, v_4\}$, and $E(G) = \{\{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_4\}, \{v_4, v_1\}\}$. Let $S_1 = \{v_1, v_2\}$, and $S_2 = \{v_1, v_3\}$. To see the unbounded price of anarchy, note that $(S_1, \bar{S}_1)$ is a cut of total happiness two and $(S_2, \bar{S}_2)$ is a Nash equilibrium of total happiness zero.

Moreover, we know that the expected total happiness of a random cut is zero. A random cut cut is a cut that is picked uniformly at random from all possible cuts. A random cut is not an approximate solution for this social function as it is a $\frac{1}{2}$-approximation for the cut value. Here, we prove that in unweighted graphs of large girth, if we start from a random cut, the total happiness of the cut after a random one-round path is an approximate solution. In fact, this gives a sub-logarithmic-approximation algorithm for the total happiness objective function.

For some $\delta > 0$, we call an edge of G, *δ-good*, if at least one of its end-points, has degree at most δ. Also, we call an edge of G, *δ-bad*, if it is not δ-good.

Lemma 1 *Let $G = (V(G), E(G))$, be a graph with $|E(G)| \leq k|V(G)|$. Then, the number of δ-good edges of G, is at least $\frac{\delta+1-2k}{\delta+1}n$.*

Proof. Since $|E(G)| \leq kn$, the average degree of G is at most $2k$. If we pick a vertex $v \in V(G)$, uniformly at random, we have $\mathbf{Pr}[\deg(v) \leq \delta] = 1 - \mathbf{Pr}[\deg(v) \geq \delta + 1] \geq 1 - \frac{2k}{\delta+1} = \frac{\delta+1-2k}{\delta+1}$. Thus, at least $\frac{\delta+1-2k}{\delta+1}n$ vertices have degree at most δ. It follows that G has at least $\frac{\delta+1-2k}{2\delta+2}$, δ-good edges. $\square$

Consider the cut $(S, \bar{S})$, at the end of a random one-round path. Let $\prec$ be the total order on the elements of $V(G)$, defined by the random ordering of the nodes in the random one-round path. For each $e \in E(G)$, let X_e be an indicator random variable, such that $X_e = 1$, if one end-point of e is in S, and the other is in $\bar{S}$, and $X_e = 0$, otherwise.

For a pair $u, v \in V(G)$, let $\mathcal{E}_{u,v}$ denote the event that there exists a path $p = x_1, x_2, \ldots, x_{|p|}$, with $u = x_1$, and $v = x_{|p|}$, and for any i, with $1 \leq i < |p|$, $x_i \prec x_{i+1}$.

Lemma 2 *Let $\{u, v\}, \{v, w\} \in E(G)$, such that $u \prec w \prec v$. Then, for any $C' > 0$, there exists a constant $C = C(C')$, such that if the girth of G is at least $C\frac{\log n}{\log \log n}$, then $\mathbf{Pr}[\mathcal{E}_{u,w}] < n^{-C'}$.*

Proof. Since $w \prec v$, it follows that if the event $\mathcal{E}_{u,w}$ happens, then there exists a path $p = x_1, x_2, \ldots, x_{|p|}$, which does not visit v, with $u = x_1$, $w = x_{|p|}$, and $x_i \prec x_{i+1}$, for any i, with $1 \leq i < n$.

Let g be the girth of G. Consider the subgraph G' of $G \setminus \{v\}$, induced by the the vertices that are at distance at most $g-3$ from v. Since the length of the shortest cycle of G is at least g, it follows that G' is a tree, and it does not contain w. Let $\mathcal{P}$ be the set of all paths that start from u, have length $g-3$, and do not visit v. Since G' is a tree, $\mathcal{P}$ contains less than n paths. Clearly, if a path p that satisfies the above conditions exists, then there is a path $p' \in \mathcal{P}$, with $p = x'_1, \ldots, x'_{g-2}$, such that for any i, with $1 \leq i < g-2$, $x'_i \prec x'_{i+1}$. Thus, for a sufficiently large constant C, the probability that such a path exists, is less than $n / \left(C \frac{\log n}{\log \log n} - 3\right)! < n^{-C'}$.
□

Lemma 3 *For any* $e \in E(G)$*, we have* $\mathbf{Pr}[X_e = 1] \geq 1/2 - o(1)$.

Proof. Let $e = \{u, v\}$, and assume w.l.o.g., that $u \prec v$. If u is the only neighbor of v, that precedes v, w.r.to $\prec$, then clearly $\mathbf{Pr}[X_e = 1] = 1$.

Assume now that there exists $u' \in V(G)$, $u' \neq u$, with $\{u', v\} \in E(G)$, and $u' \prec v$. By Lemma 2, it follows that $\mathbf{Pr}[\mathcal{E}_{u,u'} \vee \mathcal{E}_{u',u}] < 2/n^{C'}$. Observe that if none of the events $\mathcal{E}_{u,u'}$, and $\mathcal{E}_{u',u}$ happens, then the choice of the part of the cut that u belongs, is independent of the choice of the part that u' belongs. That is, the conditional probability that u and u' are both in S, or $\bar{S}$ is $1/2$.

Since v has at most n neighbors, it follows that the probability that there exists neighbors u_1, u_2 of v, such that $\mathcal{E}_{u_1,u_2}$, is at most $O(1/n)$. Thus, with probability at least $1 - O(1/n)$, none of these events happens. In this case, the conditional probability that $X_e = 1$, is at least $1/2$. It follows that $Pr[X_e = 1] \geq 1/2 - O(1/n)$.
□

Lemma 4 *Let* $e = \{u, v\} \in E(G)$*, with* $u \prec v$*, and* $\deg(v) \leq \delta$*. Then,* $\mathbf{Pr}[X_e = 1] \geq 1/2 + \Omega(1/\sqrt{\delta})$.

Proof. By applying the same argument of the proof of Lemma 3, we obtain that the probability that there exists neighbors u_1, u_2 of v, such that $\mathcal{E}_{u_1,u_2}$, is at most $O(1/n)$. Thus, with probability at least $1 - o(1)$, none of these events happens.

Assume now that none of these events happens. For each neighbor w, of v, let Y_w be an indicator random variable, such that $Y_w = 1$, if w is in the same part of the cut with u, and $Y_w = 0$, otherwise. Let $Y = \sum_{\{w,v\} \in E(G)} Y_w$. Since $\mathbf{Pr}[Y_u = 1] = 1$, we obtain $\mathbf{E}[Y] = (d+1)/2$. We will consider two cases for δ.

Case 1: If δ is odd, we have

$$\begin{aligned} \mathbf{Pr}[\{u,v\} \text{ is cut}] &\geq \mathbf{Pr}[Y \geq (\delta+1)/2] \\ &= \mathbf{Pr}[Y = (\delta+1)/2] + \mathbf{Pr}[Y > (\delta+1)/2] \end{aligned}$$

Note that

$$\begin{aligned} \mathbf{Pr}[Y = (\delta+1)/2] &= 2^{-\delta+1} \binom{\delta - 1}{\frac{\delta-1}{2}} \\ &= \Omega(1/\sqrt{\delta}) \end{aligned}$$

Since $\mathbf{Pr}[Y > (\delta+1)/2] = \mathbf{Pr}[Y < (\delta+1)/2]$, we obtain $\mathbf{Pr}[\{u,v\} \text{ is cut}] = 1/2 + \Omega(1/\sqrt{\delta})$.

Case 2: If δ is even, we have

$$\begin{aligned} \mathbf{Pr}[\{u,v\} \text{ is cut}] &= \frac{1}{2}\mathbf{Pr}[Y = \delta/2] + \mathbf{Pr}[Y > \delta/2] \\ &= \frac{1}{2}\mathbf{Pr}[Y = \delta/2] + \frac{1}{2} \end{aligned}$$

Note that

$$\begin{aligned}\mathbf{Pr}[Y = \delta/2] &= 2^{-\delta+1}\binom{\delta - 1}{\frac{\delta}{2} - 1} \\ &= \Omega(1/\sqrt{\delta})\end{aligned}$$

Thus, we obtain $\mathbf{Pr}[\{u, v\} \text{ is cut}] = 1/2 + \Omega(1/\sqrt{\delta})$.

$\square$

Theorem 15 *For any unweighted simple graph of girth at least $2\log n$, if we start from a random cut, the expected value of the happiness at the end of a random one-round path, is within a constant factor from the maximum happiness.*

Proof. If $G = (V(G), E(G))$ is a graph of girth at least $2\log n$, then $|E(G)| \leq 3n$. Also, by Lemma 1, it follows that there are at least $n/9$, 8-good edges in G. By Lemma 3, it follows that the probability that an 8-bad edge is cut, is at least $1/2 - o(1)$, while by Lemma 4, the probability that an 8-good edge is cut, is at least $1/2 + \Omega(1)$. Thus, the expectation of the total happiness after a random one-round path, is $\Omega(n)$. $\square$

Theorem 16 *There exists a constant C', such that for any $C > C'$, and for any unweighted simple graph of girth at least $C\frac{\log n}{\log\log n}$, if we start from a random cut, the expected value of the happiness at the end of a random one-round path, is within a $\frac{1}{(\log n)^{O(1/C)}}$ factor from the maximum happiness.*

Proof. We have $|E(G)| \leq n + n^{1+1/\lfloor\frac{C\log n}{2\log\log n}\rfloor} < n + n^{1+\frac{C\log n}{2\log\log n}}$, and for sufficiently large n, $|E(G)| = O(n\log^{1/C} n)$. Also, by Lemma 1, it follows that there are at least $\Omega(n)$, $\log^{1/C} n$-good edges in G. By Lemma 3, the probability that a $\log^{1/C} n$-bad edge is cut, is at least $1/2 - o(1)$, while by Lemma 4, the probability that a $\log^{1/C} n$-good edge is cut, is at least $1/2 + \Omega(\log^{-1/2C} n)$. Thus, the expectation of the total happiness after a random one-round path, is $\Omega(n\log^{-1/2C} n)$. $\square$

Note that the above theorem also gives a combinatorial sub-logarithmic approximation algorithm for the total happiness problem in unweighted graphs of large girth. As mentioned before, this objective function is considered in the context of correlation clustering problem [1] and a $\log(n)$-approximation is recently known for this function in general graphs [3].

2.6 Conclusion and Open Problems

In this paper, we give lower and upper bounds for the convergence to approximate solutions in cut games. We believe that in order to capture the computational issues of the performance of systems under lack of coordination, instead of just bounding the performance of a Nash equilibrium, it is necessary to bound the performance of the system along paths induced by a polynomial number of movements by players. We are especially interested in the performance of the system along fair paths (e.g., covering paths) to equilibria, since we may hope for better social functions along these paths. This, in turn, has implications in local search method in optimization. In order to use local search to optimize a function, we do not need to find the local optimum. We can find short paths to approximate solutions, e.g. by randomizing over the choice of the next local operation. Similar questions can be asked about different classes of games and local optimization problems.

Among the problems for the convergence in the party affiliation games, we do not know if the result of Theorem 11 for random one-round paths holds with high probability or not. The complexity of finding an approximate Nash equilibrium in the above cut game is not known to us. Another open problem is bounding the length of paths to Nash equilibria in cut games in which all players are $(1+\epsilon)$-greedy.

Acknowledgments. We would like to thank Michel Goemans for many helpful discussions.

Chapter 3

Sink Equilibria and Convergence

abstract. We introduce a new equilibrium concept, the *sink equilibrium*. Given the best response graph induced by the set of pure strategy profiles in a game, a sink equilibrium is a minimal set of strategy profiles that is closed under the best responses of players. In other words, any best response of players from a strategy profile in the sink equilibrium goes to another strategy profile in the sink equilibrium. We argue that there is a natural convergence process to sink equilibria in games where agents use pure strategies. This leads to an alternative measure of the social cost of a lack of coordination, the price of sinking. One advantage of sink equilibria is that it always exists. We illustrate the importance of this measure in two ways. First, we show that it more accurately reflects the inefficiency of uncoordinated solutions in competitive games. In particular, we give an example (a valid-utility games) in which the game converges to solutions which are a factor n worse than socially optimal. The price of sinking is at least n, but the price of anarchy for mixed Nash equilibrium is at most 2. We also show that the result is almost tight; the price of sinking is at most $n + 1$ in valid-utility games. Second, sink equilibria always exist. Thus, even in games in which pure strategy Nash equilibria do not exist, we can still calculate the price of sinking. As an example, we bound the price of sinking for a distributed caching game. Finally, we present a hardness result which shows that, in general, it is PLS-complete to find a sink equilibrium of a valid-utility game. This, in turn, implies existence of exponentially best response paths to sink equilibria in this class of games.

Bibliography

[1] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56:89–113, 2004.

[2] P. Brucker, J. Hurink, and F. Werner. Improving local search heuristics for some scheduling problems. *Discere Applied Mathematics*, 72:47–69, 1997.

[3] M. Charikar and A. Wirth. Maximizing quadratic programs: extending grothendieck's inequality. In *FOCS*, page to appear, 2004.

[4] E. Even-dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *ICALP*, pages 502–513, 2003.

[5] A. Fabrikant, C. Papadimitriou, and K. Talwar. On the complexity of pure equilibria. In *STOC*, 2004.

[6] M. Goemans, L. Li, V.S.Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *MOBIHOC*, 2004.

[7] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.

[8] D. Johnson, C.H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Computer and System Sciences*, 37:79–100, 1988.

[9] I. Milchtaich. Congestion games with player-specific. *Games and Economics Behavior*, 13, 1996.

[10] V.S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *RANDOM-APPROX*, pages 183–194, 2004.

[11] D. Monderer and L. Shapley. Potential games. *Games and Economics Behavior*, 14:124–143, 1996.

[12] N. Nisan and A. Ronen. Algorithmic mechanism design. In *STOC*, 1999.

[13] C. Papadimitriou. Algorithms, games, and the internet. In *STOC*, 2001.

[14] S. Poljak. Integer linear programs and local search for max-cut. *Siam Journal of Computing*, 24(4):822–839, 1995.

[15] T. Roghgarden and E. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economics Behavior*, 2002.

[16] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[17] T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[18] A. Schaffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.

[19] P. Schuurman and T. Vredeveld. Performance guarantees of local search for multiprocessor scheduling. In *IPCO*, pages 370–382, 2001.

[20] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of 43rd Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.