

Convergence and Approximation in Potential Games

George Christodoulou¹ *, Vahab S. Mirrokni² **, and Anastasios Sidiropoulos²

¹ National and Kapodistrian University of Athens
Dept. of Informatics and Telecommunications

² MIT Computer Science and Artificial Intelligence Laboratory

Abstract. We study the speed of convergence to approximately optimal states in two classes of potential games. We provide bounds in terms of the number of *rounds*, where a round consists of a sequence of movements, with each player appearing at least once in each round. We model the sequential interaction between players by a *best-response walk* in the state graph, where every transition in the walk corresponds to a best response of a player. Our goal is to bound the social value of the states at the end of such walks. In this paper, we focus on two classes of potential games: selfish routing games, and cut games (or party affiliation games [7]).

1 Introduction

The main tool for analyzing the performance of systems where selfish players interact without central coordination, is the notion of the *price of anarchy* in a game [16]; this is the worst case ratio between an optimal social solution and a Nash equilibrium. Intuitively, a high price of anarchy indicates that the system under consideration requires central regulation to achieve good performance. On the other hand, a low price of anarchy does not necessarily imply high performance of the system. One main reason for this phenomenon is that in many games, the repeated selfish behavior of players may not lead to a Nash equilibrium. Moreover, even if the selfish behavior of players converges to a Nash equilibrium, the *rate* of convergence might be very slow. Thus, from a practical and computational viewpoint, it is important to evaluate the rate of convergence to approximate solutions.

By modeling the repeated selfish behavior of the players as a sequence of atomic improvements, the resulting convergence question is related to the running time of local search algorithms. In fact, the theory of PLS-completeness [22]

* E-mail: gchristo@di.uoa.gr. Research supported in part by the programme AEOLUS FP6-015964. Research supported in part by the programme pythagoras II, which is funded by the European Social Fund (75%) and the Greek Ministry of Education (25%).

** E-mail: mirrokni@theory.csail.mit.edu. Part of this work was done while the author was with the Strategic Planning and Optimization Group at Amazon.com.

*** E-mail: tasos@theory.csail.mit.edu

and the existence of exponentially long walks in local optimization problems such as Max-2SAT and Max-Cut, indicate that in many of these settings, we cannot hope for a polynomial-time convergence to a Nash equilibrium. Therefore, for such games, it is not sufficient to just study the value of the social function at Nash equilibria. To deal with this issue, we need to bound the social value of a strategy profile after *polynomially many* best-response improvements by players.

Potential games are games in which any sequence of improvements by players converges to a pure Nash equilibrium. Equivalently, in potential games, there is no cycle of strict improvements of players. This is equivalent to the existence of a potential function that is strictly increasing after any strict improvement. In this paper, we study the speed of convergence to approximate solutions in two classes of potential games: selfish routing (or congestion) games and cut games.

Related Work. This work is motivated by the negative results of the convergence in congestion games [7], and the study of convergence to approximate solutions games [14, 11]. Fabrikant, Papadimitriou, and Talwar [7] show that for general congestion and asymmetric selfish routing games, the problem of finding a pure Nash equilibrium is PLS-complete. This implies exponentially long walks to equilibria for these games. Our model is based on the model introduced by Mirrokni and Vetta [14] who addressed the convergence to approximate solutions in basic-utility and valid-utility games. They prove that starting from any state, one round of selfish behavior of players converges to a $1/3$ -approximate solution in basic-utility games. Goemans, Mirrokni, and Vetta [11] study a new equilibrium concept (i.e. sink equilibria) inspired from convergence on best-response walks and proved fast convergence to approximate solutions on random best-response walks in (weighted) congestion games. In particular, their result on the price of sinking of the congestion games implies polynomial convergence to constant-factor solutions on random best-response walks in selfish routing games with linear latency functions. Other related papers studied convergence for different classes of games such as load balancing games [6], market sharing games [10], and distributed caching games [8].

A main subclass of potential games is the class of *congestion games* introduced by Rosenthal [18]. Monderer and Shapley [15] proved that congestion games are equivalent to the class of *exact potential games*. In an exact potential game, the increase in the payoff of a player is equal to the increase in the potential function. Both selfish routing games and cut games are a subclass of exact potential games, or equivalently, congestion games. Tight bounds for the price of anarchy is known for both of these games in different settings [19, 1, 5, 4]. Despite all the recent progress in bounding the price of anarchy in these games, many problems about the speed of convergence to approximate solutions for them are still open.

Two main known results for the convergence of selfish routing games are the existence of exponentially long best-response walks to equilibria [7] and fast convergence to constant-factor solutions on random best-response walks [11]. To the best of our knowledge, no results are known for the speed of convergence to approximate solutions on deterministic best-response walks in the general selfish routing game. Preliminary results of this type in some special load balancing

games are due to Suri, Tóth and Zhou [20, 21]. Our results for general selfish routing games generalize their results.

The Max-Cut problem has been studied extensively [12], even in the local search setting. It is well known that finding a local optimum for Max-Cut is PLS-complete [13, 22], and there are some configurations from which walks to a local optimum are exponentially long. In the positive side, Poljak [17] proved that for cubic graphs the convergence to a local optimum requires at most $O(n^2)$ steps. The total happiness social function is considered in the context of correlation clustering [2], and is similar to the total agreement minus disagreement in that context. The best approximation algorithm known for this problem gives a $O(\log n)$ -approximation [3], and is based on a semidefinite relaxation.

Our Contribution. Our work deviates from bounding the distance to a Nash equilibrium [22, 7], and focuses in studying the rate of convergence to an approximate solution [14, 11]. We consider two types of walks of best responses: random walks and deterministic fair walks. On random walks, we choose a random player at each step. On deterministic fair walks, the time complexity of a game is measured in terms of the number of *rounds*, where a round consists of a sequence of movements, with each player appearing at least once in each round.

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

We also study the convergence in *cut games*, that are motivated by the *party affiliation game* [7], and are closely related to the local search algorithm for the Max-Cut problem [22]. 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, \dots, s_n) is $\sum_j s_j s_i w_{ij}$. The weight of an edge corresponds to the level of *disagreement* of the endpoints of that edge. This game models the clustering of a society into two parties that minimizes the disagreement within each party, or maximizes the disagreement between different parties. Such problems play a key role in the study of social networks.

We can model the party affiliation game as the following cut game: each vertex of a graph is a player, with payoff its contribution in the cut (i.e. the total weight of its adjacent edges that have endpoints in different parts of the cut). It follows that a player moves if he can improve his contribution in the cut, or equivalently, he can improve the value of the cut. The pure Nash equilibria exist in this game, and selfish behavior of players converges to a Nash equilibrium.

We consider two social functions: the cut and the total happiness, defined as the value of the cut minus the weight of the rest of edges. First, we prove *fast convergence on random walks*. More precisely, the selfish behavior of players in a round in which the ordering of the player is picked uniformly at random, results

in a cut that is a $\frac{1}{8}$ -approximation in expectation. We complement our positive results by examples that exhibit *poor deterministic convergence*. That is, we show the existence of fair walks with *exponential* length, that result in a poor social value. We also model the selfish behavior of *mildly greedy* players that move if their payoff increases by at least a factor of $1 + \epsilon$. We prove that in contrast to the case of (totally) greedy players, mildly greedy players converge to a constant-factor cut after one round, under any ordering. For unweighted graphs, we give an $\Omega(\sqrt{n})$ lower bound and an $O(n)$ upper bound for the number of rounds required in the worst case to converge to a constant-factor cut.

Finally, for the total happiness social function, we show that for unweighted graphs of large girth, starting from a random configuration, greedy behavior of players in a random order converges to an approximate solution after one round. We remark that this implies a combinatorial algorithm with sub-logarithmic approximation ratio, for graphs of sufficiently large girth, while the best known approximation ratio for the general problem is $O(\log n)$ [3], and is obtained using semidefinite programming.

2 Definitions and Preliminaries

In order to model the selfish behavior of players, we use the notion of a *state graph*. Each vertex in the state graph represents a *strategy state* $S = (s_1, s_2, \dots, s_n)$, and corresponds to a pure strategy profile (e.g an allocation for a congestion game, or a cut for a cut game). The arcs in the state graph correspond to best response moves by the players.

Definition 1. A state graph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is a directed graph, where each vertex in \mathcal{V} corresponds to a strategy state. There is an arc from state S to state S' with label j iff by letting player j play his best response in state S , the resulting state is S' .

Observe that the state graph may contain loops. A *best response walk* is a directed walk in the state graph. We say that player i plays in the best response walk \mathcal{P} , if at least one of the edges of \mathcal{P} has label i . Note that players play their best responses sequentially, and not in parallel. Given a best response walk starting from an arbitrary state, we are interested in the social value of the last state on the walk. Notice that if we do not allow every player to make a best response on a walk \mathcal{P} , then we cannot bound the social value of the final 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 value³. Motivated by this simple observation, we introduce the following models that capture the intuitive notion of a fair sequence of moves.

One-round walk: Consider an arbitrary ordering of all players i_1, \dots, i_n . A walk \mathcal{P} of length n in the state graph is a *one-round walk* if for each $j \in [n]$, the j th edge of \mathcal{P} has label i_j .

³ e.g. in the cut social function, most of the weight of the edges of the graph might be concentrated to the edges that are adjacent to a single vertex.

Covering walk: A walk \mathcal{P} in the state graph is a *covering walk* if for each player i , there exists an edge of \mathcal{P} with label i .

k -Covering walk: A walk \mathcal{P} in the state graph is a *k -covering walk* if there are k covering walks $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k$, such that $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k)$.

Random walk: A walk \mathcal{P} in the state graph is a *random walk*, if at each step the next player is chosen uniformly at random.

Random one-round walk: Let σ be an ordering of players picked uniformly at random from the set of all possible orderings. Then, the one-round walk \mathcal{P} corresponding to the ordering σ , is a *random one-round walk*.

Note that unless otherwise stated, all walks are assumed to start from an arbitrary initial state. This model has been used by Mirrokni and Vetta [14], in the context of extensive games with complete information.

Congestion games. A congestion game is defined by a tuple $(N, E, (\mathcal{S}_i)_{i \in N}, (f_e)_{e \in E})$ where N is a set of players, E is a set of facilities, $\mathcal{S}_i \subseteq 2^E$ is the pure strategy set for player i : a pure strategy $s_i \in \mathcal{S}_i$ for player i is a set of facilities, and f_e is a latency function for the facility e depending on its load. We focus on linear delay functions with nonnegative coefficients; $f_e(x) = a_e x + b_e$.

Let $S = (s_1, \dots, s_N) \in \times_{i \in N} \mathcal{S}_i$ be a state (strategy profile) for a set of N players. The cost of player i , in a state S is $c_i(S) = \sum_{e \in s_i} f_e(n_e(S))$ where by $n_e(S)$ we denote the number of players that use facility e in S . The objective of a player is to minimize its own cost. We consider as a social cost of a state S , the sum of the players' costs and we denote it by $C(S) = \sum_{i \in N} c_i(S) = \sum_{e \in E} n_e(S) f_e(n_e(S))$.

In weighted congestion games, player i has weighted demand w_i . By $\theta_e(S)$, we denote the total load on a facility e in a state S . The cost of a player in a state S is $c_i(S) = \sum_{e \in s_i} f_e(\theta_e(S))$. We consider as a social cost of a state S , the weighted sum $C(S) = \sum_{i \in N} w_i c_i(S) = \sum_{e \in E} \theta_e(S) f_e(\theta_e(S))$. We will use subscripts to distinguish players and superscripts to distinguish states.

Note that the selfish routing game is a special case of congestion games. Although we state all the results for congestion games with linear latency functions, all of the results (including the lower and upper bounds) hold for selfish routing games.

Cut Games. In a cut game, we are given an undirected graph $G(V, E)$, with n vertices and edge weights $w : E(G) \rightarrow \mathbb{Q}^+$. We will always assume that G is connected, simple, and does not contain loops. For each $v \in V(G)$, let $\deg(v)$ be the degree of v , and let $\text{Adj}(v)$ be the set of neighbors of v . Let also $w_v = \sum_{u \in \text{Adj}(v)} w_{uv}$. A cut in G is a partition of $V(G)$ into two sets, T and $\bar{T} = V(G) - T$, and is denoted by (T, \bar{T}) . The value of a cut is the sum of edges between the two sets T and \bar{T} , i.e. $\sum_{v \in T, u \in \bar{T}} w_{uv}$.

The *cut game* on a graph $G(V, E)$, is defined as follows: each vertex $v \in V(G)$ is a player, and the strategy of v is to chose one side of the cut, i.e. v can chose $s_v = -1$ or $s_v = 1$. A strategy profile $S = (s_1, s_2, \dots, s_n)$, corresponds to a cut (T, \bar{T}) , where $T = \{i | s_i = 1\}$. The payoff of player v in a strategy profile S , denoted by $\alpha_v(S)$, is equal to the contribution of v in the cut, i.e.

$\alpha_v(S) = \sum_{i:s_i \neq s_v} w_{iv}$. It follows that the cut value is equal to $\frac{1}{2} \sum_{v \in V} \alpha_v(S)$. If S is clear from the context, we use α_v instead of $\alpha_v(S)$ to denote the payoff of v . We denote the maximum value of a cut in G , by $c(G)$. The *happiness* of a vertex v is equal to $\sum_{i:s_i \neq s_v} w_{iv} - \sum_{i:s_i = s_v} w_{iv}$.

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 the payoffs of vertices. The second social function is half the sum of the happiness of vertices. We call the second social function, *total happiness*.

3 Congestion Games

In this section, we focus on the convergence to approximate solutions in congestion games with linear latency functions. It is known [15, 18] that any best-response walk on the state graph leads to a pure Nash equilibrium, and a pure equilibrium is a constant-factor approximate solution [1, 5, 4]. Unless otherwise stated, we assume without loss of generality, that the players' ordering is $1, \dots, N$.

3.1 Upper Bounds for One-round Walks

In this section, we bound the total delay after one round of best responses of players. We prove that starting from an arbitrary state, the solution after one round of best responses is a $\Theta(N)$ -approximate solution. We will also prove that starting from an *empty* state, the approximation factor after one round of best responses is a constant factor. This shows that the assumption about the initial state is critical for this problem.

Theorem 1. *Starting from an arbitrary initial state S^0 , any one-round walk \mathcal{P} leads to a state S^N that has approximation ratio $O(N)$.*

Proof. Let X be the optimal allocation and $S^i = (s_1^i, \dots, s_i^i, s_{i+1}^0, \dots, s_N^0)$ an intermediate state. Let $m_e(S^i), k_e(S^i)$ be the number of the players of the final and of the initial state respectively, using facility e in a state S^i , and $M(S^i), K(S^i)$ the corresponding sums. Clearly $n_e(S^i) = m_e(S^i) + k_e(S^i)$ and $K(S^i) = K(S^{i-1}) - \sum_{e \in s_i^0} (a_e - b_e - 2a_e k_e(S^{i-1}))$. By summing over all intermediate states and using the fact $K(S^N) = 0$, it follows that:

$$K(S^0) = C(S^0) = \sum_{e \in E} k_e(S^0) f_e(k_e(S^0)) = \sum_{i \in N} \sum_{e \in s_i^0} (2a_e k_e(S^{i-1}) - a_e + b_e) \quad (1)$$

Since player i in state S^{i-1} prefers strategy s_i^N than x_i , we get

$$\sum_{e \in s_i^N} f_e(n_e(S^{i-1})) + \sum_{e \in s_i^N - s_i^0} a_e \leq \sum_{e \in x_i} f_e(n_e(S^{i-1}) + 1)$$

For every intermediate state S^i , the social cost is

$$C(S^i) = C(S^{i-1}) + \sum_{e \in s_i^N - s_i^0} (2a_e n_e(S^{i-1}) + a_e + b_e) + \sum_{e \in s_i^0 - s_i^N} (a_e - b_e - 2a_e n_e(S^{i-1}))$$

Summing over all intermediate states and using equality (1), we get

$$\begin{aligned}
C(S^N) &= \sum_{i \in N} \sum_{e \in s_i^N - s_i^0} (2a_e n_e(S^{i-1}) + a_e + b_e) + \sum_{i \in N} \sum_{e \in s_i^0} (2a_e k_e(S^{i-1}) - a_e + b_e) \\
&\quad + \sum_{i \in N} \sum_{e \in s_i^0 - s_i^N} (a_e - b_e - 2a_e n_e(S^{i-1})) \\
&= \sum_{i \in N} \sum_{e \in s_i^N - s_i^0} (2a_e n_e(S^{i-1}) + a_e + b_e) + \sum_{i \in N} \sum_{e \in s_i^0 \cap s_i^N} (2a_e k_e(S^{i-1}) - a_e + b_e) \\
&\quad - \sum_{i \in N} \sum_{e \in s_i^0 - s_i^N} 2a_e m_e(S^{i-1}) \\
&\leq 2 \sum_{i \in N} \sum_{e \in s_i^N} f_e(n_e(S^{i-1})) + 2 \sum_{i \in N} \sum_{e \in s_i^N - s_i^0} a_e \\
&\leq \sum_{i \in N} \sum_{e \in x_i} 2f_e(n_e(S^{i-1}) + 1) \\
&\leq \sum_{i \in N} \sum_{e \in x_i} 2f_e(N + 1) = \sum_{e \in E} 2n_e(X) f_e(N + 1) = O(N)C(X)
\end{aligned}$$

□

In the next section, we will show that the above bound is tight up to a constant factor. As mentioned earlier, the assumption about the initial state is critical for this problem. We will call a state *empty*, if no player is committed to any of its strategies. Note that the one-round walk starting from an empty state is essentially equivalent to the greedy algorithm for a generalized scheduling problem, where a task may be assigned into many machines. Suri et al. [20, 21] address similar questions for the special case of the congestion games where the available strategies are single sets (i.e. each player can choose just one facility). They give a 3.08 lower bound and a 17/3 upper bound. For the special case of identical facilities (equal speed machines) they give an upper bound of $\frac{(\phi+1)^2}{\phi} \approx 4.24$. We generalize this result for our more general setting.

Theorem 2. *Starting from the empty state S^0 , any one-round walk \mathcal{P} leads to a state S^N that has approximation ratio of at most $\frac{(\phi+1)^2}{\phi} \approx 4.24$.*

Now we turn our attention to weighted congestion games with linear latency functions, where player i has weighted demand w_i . Fotakis et al. [9] showed that this game with linear latency functions is a potential game.

Theorem 3. *In weighted congestion games with linear latency functions, starting from the initial empty state S^0 , any one-round walk \mathcal{P} leads to a state S^N that has approximation ratio of at most $(1 + \sqrt{3})^2 \approx 7.46$.*

3.2 Lower Bounds

The next theorem shows that the result of Theorem 1 is tight and explains why it is necessary in the upper bounds given above to consider walks starting from an empty allocation.

Theorem 4. *For any $N > 0$, there exists an N -player instance of the unweighted congestion game, and an initial state S^0 such that for any one-round walk \mathcal{P} starting from S^0 , the state at the end of \mathcal{P} is an $\Omega(N)$ -approximate solution.*

Proof. Consider $2N$ players and $2N + 2$ facilities $\{0, 1, \dots, 2N + 1\}$. The available strategies for the first players are $\{\{0\}, \{i\}, \{N + 1, \dots, 2N\}\}$ and for the N last $\{\{2N + 1\}, \{i\}, \{1, \dots, N\}\}$. In the initial allocation, every player plays its third strategy. Consider any order on the players and let them begin to choose their best responses. It is easy to see that in the first steps, the players would prefer their first strategy. If this happens until the end of the round, the resulting cost is $\Omega(N^2)$. Thus, we can assume that at some step, the $(k + 1)$ -th player from the set $\{1, \dots, N\}$ prefers his second strategy while all the previous k players of the same set have chosen their first strategies. The status of the game at this step is as follows: k players of the first group play their first strategy, m players of the second group play their first strategy and the remaining players play their initial strategy. Since player $k + 1$ prefers his second strategy, this means $k = N - m$ and so one of the m, N is at least $N/2$. The cost at the end will be at least $m^2 + k^2 + N = \Omega(N^2)$. On the other hand, in the optimal allocation everybody chooses its second strategy which gives cost $2N$. Thus, the approximation ratio is $\Omega(N)$. \square

We extend theorem 4 for the case of t -covering walks, for $t > 1$. We remark that the following result holds only for a fixed ordering of the players.

Theorem 5. *For any $t > 0$, and for any sufficiently large $N > 0$, there exists an N -player instance of the unweighted congestion game, an initial state S^0 , and an ordering σ of the players, such that starting from S^0 , after t rounds where the players play according to σ , the cost of the resulting allocation is a $(N/t)^\epsilon$ -approximation, where $\epsilon = 2^{-O(t)}$.*

4 Cut Games: The Cut Social Function

4.1 Fast Convergence on Random Walks

First we prove positive results for the convergence to constant-factor approximate solutions with random walks. We show that the expected value of the cut after a random one-round walk is within a constant factor of the maximum cut.

Theorem 6. *In weighted graphs, the expected value of the cut at the end of a random one-round walk is at least $\frac{1}{8}$ of the maximum cut.*

Proof. It suffices to show that after a random one-round walk, for every $v \in V(G)$, $\mathbf{E}[\alpha_v] \geq \frac{1}{8}w_v$.

Consider a vertex v . The probability that v occurs after exactly k of its neighbors is $\frac{1}{\deg(v)+1}$. After v moves, the contribution of v in the cut is at least $\frac{w_v}{2}$. Conditioning on the fact 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 at most w_{uv} . Thus the expected 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 obtain $\mathbf{E}[\alpha_v] \geq \sum_{k=0}^{\deg(v)} \frac{1}{\deg(v)+1} \max(0, w_v(\frac{1}{2} - \frac{\deg(v)-k}{\deg(v)})) = \frac{w_v}{\deg(v)+1} \sum_{k=0}^{\lfloor \frac{\deg(v)}{2} \rfloor + 1} \frac{2k - \deg(v)}{2\deg(v)} \geq \frac{w_v}{8}$. The result follows by the linearity of expectation. \square

The next theorem studies a random walk of best responses (not necessarily a one-round walk).

Theorem 7. *There exists a constant $c > 0$ such that the expected value of the cut at the end of a random walk of length $cn \log n$ is a constant-factor of the maximum cut.*

4.2 Poor Deterministic Convergence

We now give lower bounds for the convergence to approximate solutions for the cut social function. First, we give a simple example for which we need at least $\Omega(n)$ rounds of best responses to converge to a constant-factor cut. The construction resembles a result of Poljak [17].

Theorem 8. *There exists a weighted graph $G(V, E)$, 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.*

We next combine a modified version of the above construction with a result of Schaffer and Yannakakis for the Max-Cut local search problem [22], to obtain an exponentially-long walk with poor cut value.

Theorem 9. *There exists a weighted graph $G(V, E)$, with $|V(G)| = \Theta(n)$, and a k -covering walk \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 [22], it is shown that there exists a weighted graph $G_0(V, E)$, and an initial cut (T_0, \bar{T}_0) , such that the length of *any* walk in the state graph, from (T_0, \bar{T}_0) to a pure strategy Nash equilibrium, is exponentially long. Consider such a graph of size $\Theta(n)$, with $V(G_0) = \{v_0, v_1, \dots, v_N\}$. Let \mathcal{P}_0 be an exponentially long walk from (T_0, \bar{T}_0) to a Nash equilibrium in which we let vertices v_0, v_1, \dots, v_N play in this order for exponential an number of rounds. Let $S_0, S_1, \dots, S_{|\mathcal{P}_0|}$ be the sequence of states visited by \mathcal{P}_0 and let y_i be the vertex that plays his best response from state S_i to state S_{i+1} . The result of [22] guarantees that there exists a vertex, say v_0 , which wants to change side (i.e. strategy) an exponential number of times along the walk \mathcal{P}_0 (since otherwise we can find a small walk to a pure Nash equilibrium). Let $t_0 = 0$, and for $i \geq 1$, let t_i be the time in which v_0 changes side for the i -th time along the walk \mathcal{P}_0 . For $i \geq 1$, let \mathcal{Q}_i be the sequence of vertices $y_{t_{i-1}+1}, \dots, y_{t_i}$. Observe that each \mathcal{Q}_i contains all of the vertices in G_0 .

Consider now a graph G , which consists of a path $L = x_1, x_2, \dots, x_n$, and a copy of G_0 . For each $i \in \{1, \dots, n-1\}$, 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 $i \in \{1, \dots, n\}$, 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 vertices in G_0 , are independent of the positions of the vertices of the path L in the current cut.

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

We are now ready to describe the exponentially long path in the state graph. Assume w.l.o.g., that in the initial cut for G_0 , we have $v_0 \in T_0$. The initial cut for G is (T, \bar{T}) , with $T = \{x_1\} \cup T_0$, and $\bar{T} = \{x_2, \dots, x_n\} \cup \bar{T}_0$. It is now straightforward to verify that there exists an exponentially large k , such that for any i , with $1 \leq i \leq k$, if we let the vertices of G play according to the sequence $\mathcal{Q}_1, \mathcal{R}_1, \mathcal{Q}_2, \mathcal{R}_2, \dots, \mathcal{Q}_i, \mathcal{R}_i$, then we have (see Figure 1):




Fig. 1. The cut (T_i, \bar{T}_i) along the walk of the proof of Theorem 9.

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

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 optimal cut. The result follows since each walk in the state graph induced by the sequence \mathcal{Q}_i and \mathcal{R}_i is a covering walk. □

4.3 Mildly Greedy Players

By Theorem 6, it follows that for any graph, and starting from an arbitrary cut, there exists a walk of length at most n to an $\Omega(1)$ -approximate cut, while Theorems 8 and 9, show that there are cases where a deterministic ordering of players may result to very long walks that do reach an approximately good cut.

We observe that if we change the game by assuming that a vertex changes side in the cut if his payoff is multiplied by at least a factor $1 + \epsilon$, for a constant $\epsilon > 0$, then the convergence is faster. We call such vertices $(1 + \epsilon)$ -greedy. In the following, we prove that if all vertices are $(1 + \epsilon)$ -greedy for a constant $\epsilon > 0$, then the value of the cut after any one-round walk is within a constant factor of the optimum.

Theorem 10. *If all vertices are $(1 + \epsilon)$ -greedy, then the cut value at the end of any one-round walk is within a $\min\{\frac{1}{4+2\epsilon}, \frac{\epsilon}{4+2\epsilon}\}$ factor of the optimal cut.*

4.4 Unweighted Graphs

In unweighted simple graphs, it is straight-forward to verify that the value of the cut at the end of an n^2 -covering walk is at least $\frac{1}{2}$ of the optimum. The following theorem shows that in unweighted graphs, the value of the cut after any $\Omega(n)$ -covering walk is a constant-factor approximation.

Theorem 11. *For unweighted graphs, the value of the cut after an $\Omega(n)$ -covering walk is within a constant-factor of the maximum cut.*

Proof. Consider a k -covering walk $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_k)$, where each \mathcal{P}_i is a covering walk. 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 the maximum size of a cut is at most $|E(G)|$, 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 vertices that change their side in the cut on the walk \mathcal{P}_i , and $V_2 = V(G) \setminus V_1$. Observe that when a vertex changes its side in the cut, the size of the cut increases by at least 1. Thus, $|V_1| < \frac{|E(G)|}{10n}$, and since the degree of each vertex is at most $n - 1$, it follows that the number of edges that are incident to vertices in V_1 , is less than $\frac{|E(G)|}{10}$.

On the other hand, if a vertex 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 vertex in V_2 , were in the cut, at some point during walk \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 walk \mathcal{P}_i , is at least $\frac{|E(G)| - |E(G)|/10}{2} - \frac{|E(G)|}{10} = \frac{7|E(G)|}{20}$. Since the maximum size of a cut is at most $|E(G)|$, we obtain that at the end of \mathcal{P}_i , the value of the cut is within a constant factor of the optimum. \square

Theorem 12. *There exists an unweighted graph $G(V, E)$, with $|V(G)| = n$, and an ordering of the 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.*

5 The Total Happiness Social Function

Due to space limitations, this section has been left to the full version.

Acknowledgments. We wish to thank Michel Goemans and Piotr Indyk for many helpful discussions. We also thank Ioannis Caragiannis for reading an earlier version of this work, and for improving the bound given in Theorem 7.

References

1. B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *STOC*, pages 57–66, 2005.
2. N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56:89–113, 2004.
3. M. Charikar and A. Wirth. Maximizing quadratic programs: extending Grothendieck’s inequality. In *FOCS*, page to appear, 2004.
4. G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *ESA*, 2005.
5. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC*, pages 67–73, 2005.
6. E. Even-dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibria. In *ICALP*, pages 502–513, 2003.
7. A. Fabrikant, C. Papadimitriou, and K. Talwar. On the complexity of pure equilibria. In *STOC*, 2004.
8. L. Fliescher, M. Goemans, V. S. Mirrokni, and M. Sviridenko. (almost) tight approximation algorithms for maximizing general assignment problems. 2004. Submitted.
9. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flow. In *ICALP*, 2004.
10. M. Goemans, L. Li, V.S.Mirrokn, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *MOBIHOC*, 2004.
11. M. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS*, 2005.
12. M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.
13. D. Johnson, C.H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Computer and System Sciences*, 37:79–100, 1988.
14. V.S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *RANDOM-APPROX*, pages 183–194, 2004.
15. D. Monderer and L. Shapley. Potential games. *Games and Economics Behavior*, 14:124–143, 1996.
16. C. Papadimitriou. Algorithms, games, and the Internet. In *STOC*, 2001.
17. S. Poljak. Integer linear programs and local search for max-cut. *Siam Journal of Computing*, 24(4):822–839, 1995.
18. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
19. T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
20. C. D. Tóth S. Suri and Y. Zhou. Selfish load balancing and atomic congestion games. In *SPAA*, pages 188–195, 2004.
21. C. D. Tóth S. Suri and Y. Zhou. Uncoordinated load balancing and congestion games in p2p systems. In *IPTPS*, pages 123–130, 2004.
22. A. Schaffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.