Fast Convergence to Nearly Optimal Solutions in Potential Games

Baruch Awerbuch* Yossi Azar† Amir Epstein‡ Vahab S. Mirrokni §
Alexander Skopalik ¶

Abstract

We study the speed of convergence of decentralized dynamics to approximately optimal solutions in potential games. We consider α-Nash dynamics in which a player makes a move if the improvement in his payoff is more than an α factor of his own payoff. Despite the known polynomial convergence of α-Nash dynamics to approximate Nash equilibria in symmetric congestion games [7], it has been shown that the convergence time to approximate Nash equilibria in asymmetric congestion games is exponential [23]. In contrast to this negative result, and as the main result of this paper, we show that for asymmetric congestion games with delay functions that satisfy a "bounded jump" condition, the convergence time of α-Nash dynamics to an approximate optimal solution is polynomial in the number of players, with approximation ratio that is arbitrarily close to the price of anarchy of the game. In particular, we show this polynomial convergence under the minimal liveness assumption that each player gets at least one chance to move in every T steps. We also prove that the same polynomial convergence result does not hold for (exact) best-response dynamics, showing the α-Nash dynamics is required. We extend these results for congestion games to other potential games including weighted congestion games with linear delay functions, cut games (also called party affiliation games) and market sharing games as follows.

*Dept. of Computer Science, Johns Hopkins University. E-Mail: baruch@cs.jhu.edu. Research supported by NSF grants ANIR-0240551 and CCR-0311795.

†School of Computer Science, Tel-Aviv University. E-Mail: azar@tau.ac.il. Research supported in part by the German-Israeli Foundation.

‡School of Computer Science, Tel Aviv University. E-Mail: amirep@tau.ac.il. Research supported in part by the German-Israeli Foundation.

§Theory Group, Microsoft Research, E-Mail: mirrokni@microsoft.com.

¶Dept. of Computer Science, RWTH Aachen, E-Mail: skopalik@cs.rwth-aachen.de.

1 Introduction

Computational game theory has lead already to many important insights for understanding Nash equilibria in systems under the control of self-interested agents. Prominent results for the quality of Nash equilibria include bounds on the price of anarchy, which is the ratio between the worst Nash equilibrium and the global optimal solution [21, 10, 24, 19], and for computational complexity [12, 11, 6]. Intuitively, a high price of anarchy a system indicates that it requires a coordination mechanism to achieve good performance. On the other hand, low price of anarchy does not necessarily imply good performance of the system [18, 15]. One main reason for this phenomenon is that in many games with selfish players acting in a decentralized fashion, the repeated selfish behavior of the players may not lead to a Nash equilibrium [15]. Moreover, the convergence rate might be very slow [12]. This motivates the question of whether selfish players acting in a decentralized fashion, converge to approximate solutions in a reasonable amount of time [18, 15, 8, 5].

In this paper, we address this question for the general class of congestion games, which are used to model routing, network design and other resource sharing scenarios in distributed systems [21, 2, 14]. We also consider other potential games. In a congestion game there are n players and a set of resources. The strategy of a player consists of a subset of these resources. Each resource possesses a delay function d_e, which depends on the number of players using this resource and the delay(cost) of each player is the sum of the delays associated with his selected resources.

Rosenthal [20] prove that every congestion game has a pure Nash equilibrium, by showing a potential function that is strictly decreasing after any strict improvement of a player. Thus, this property, shows that the natural "Nash Dynamics", in which players iteratively play best response converges to a pure Nash Equilibrium. It has been shown that the problem of finding pure Nash equilibria in congestion games is PLS-complete [12] even with linear latency functions [1]. This result holds even for symmetric congestion games. These results imply examples of congestion games and initial states from which in the Nash dynamics all Nash equilibria have distance exponential in the number of players n. For this reason, Chien and Sinclair [7] study convergence to approximate equilibria in symmetric congestion games. They consider α-Nash equilibria which are states in which no player can decrease his cost by more than a factor of $1 - \alpha$ by unilaterally changing his strategy. They also investigate α-Nash dynamics, in which we only allow moves that improve the cost of a player by a factor of more than $1 - \alpha$. For symmetric congestion games where each resource delay satisfies the "bounded jump assumption", they show that convergence to α-Nash equilibria occurs within a number of steps that is polynomial in the number of players [7]. Recently, Skopalik and Vöcking [23] show examples of asymmetric congestion games with n players and $O(n)$ resources and bounded jump delay functions such that there are states that have distance exponential in the number of players n to all α-Nash equilibria. Thus, the results for convergence to α-Nash equilibria appear in [7] cannot be extended to asymmetric congestion games. These negative results motivate the study of convergence to approximate solutions in asymmetric congestion games.

Mirrokni and Vetta [18] and Goemans et. al [15] study convergence to approximate solutions in load balancing games, valid-utility games, and congestion games. Christodoulou et al. [8] study the speed of convergence to approximate solutions in potential games. They show that after a constant number of rounds of α-Nash dynamics the approximation factor of the solution might be a superconstant. They also show that the approximation factor of a state after one round of Nash dynamics is $\Theta(n)$.

Our Results. In this paper, we study the convergence of unrestricted α-Nash dynamics to an approximately optimal solution in different classes of asymmetric congestion games and other potential games. We consider the unrestricted α-Nash dynamics with a *liveness* property that no player is prevented from moving for arbitrarily many steps. We consider asymmetric congestion games with resources satisfying a bounded jump condition. For $\gamma \geq 1$, a resource e satisfies the *γ-bounded jump* condition, if its delay function satisfies $d_e(t+1) \leq \gamma d_e(t)$ for all $t \geq 1$. This condition is rather weak. In particular, a resource with $d_e(t) = \gamma^t$ satisfy the γ-bounded jump condition.

We show that for asymmetric congestion games with the bounded jump condition, the unrestricted α-Nash dynamics with a liveness property converges to approximate solutions with approximation ratio of arbitrarily close to the price of anarchy in time that is polynomial in the number of players (For details, see Theorem 3.3 and Remark 3.5). This result implies fast convergence to good approximate solutions for the interesting case of polynomial latency functions of degree d. These results are in contrast to the negative results that appear in [12, 1, 23]. We also prove that the same polynomial convergence result does not hold for (exact) best-response dynamics, showing the α-Nash dynamics is required 3.4. We extend this result for other potential games. We first extend this result to weighted congestion games with linear delay functions for which we show that any unrestricted α-Nash dynamics satisfying the liveness property converges to a $(2.618 + \epsilon)$-approximate solution after polynomial number of α-moves. Furthermore, we extend the results to profit maximizing potential games including cut games (also called party affiliation games) and market sharing games. In these games, players maximize their payoff instead of minimizing their cost. For these games, we need to assume that players play a best-response α-moves, i.e., an α-move that has the maximum possible payoff. For both of these games, we show that any unrestricted α-Nash best-response dynamics satisfying the liveness property converges to a $(2 + \epsilon)$-approximate solution after polynomial number of α-moves. This is in contrast to the negative result of Christodoulou et al [8] for cut games that shows that convergence time of (exact) best-response dynamics to a constant-factor solution in this game is exponential.

Related Work. The study of convergence of Nash dynamics is related to local search problems, and PLS-complete problems introduced by Johnson et. al [16]. Fabrikant et al [12] proved that finding a pure Nash equilibrium of network congestion games is PLS-complete. Ackermann et al [1] showed that the same problem for network congestion games with linear latency functions is PLS-complete as well. Skopalik and Vöcking [23] showed that finding an approximate Nash equilibrium in congestion games is also PLS-complete.

Mirrokni and Vetta [18] initiated the study of convergence to approximate solutions in the context of load balancing games and valid-utility games [24]. They consider *covering walks* of best responses in which each player has at least one chance to play in each round. Motivated by studying the Nash dynamics and convergence to approximate solutions, Goemans et al [15] introduced sink equilibria, and proved that in weighted congestion games, random Nash dynamics converges to a constant-factor approximately optimal solution in expected polynomial time. However, they do not provide any bound for the convergence time of deterministic unrestricted Nash dynamics. In fact, in Theorem 3.4, we show a lower bound for deterministic Nash dynamics for these games, showing that the above result only holds for random Nash dynamics. Christodoulou et al [8] showed a tight bound of $\Theta(n)$ for the approximation factor of the solution after one round of α-Nash dynamics in congestion games with linear latency functions. They also showed that for congestion games with linear latency functions, after a constant rounds of Nash dynamics, players may not converge to an approximate solution. Here, we show that after a polynomial rounds of α-Nash dynamics, players converge to a constant-factor solution. Chekuri et al [5] and Charikar et al [4] studied convergence of Nash dynamics to approximate solutions in network cost sharing games.

The study of α-moves for convergence to approximate solutions has been also considered by Christoulou [8] in the context of cut games. They show that for any constant α (and not for an $\alpha = o(1)$) after one round of α-moves of players in a cut game, the value of the cut is a constant-factor approximate solution. Their proof does not handle the convergence of unrestricted dynamics. For a more complete list of results in these areas, see Mirrokni [17].

Cut Games (or party affiliation games) are potential games defined on an edge-weighted graph [12, 22, 8]. Nash dynamics for these games correspond to the local search algorithm for the Max-Cut problem. Schaffer and Yannakakis [22] proved that finding a Nash equilibrium in this game is PLS-complete. Christodoulou et. al [8] showed an exponential lower bound for the convergence time of (exact) best-response dynamics to constant-factor approximate solutions in these games. In contrast, we show polynomial convergence of α-Nash best-response dynamics in these games. Market sharing games are a special case of profit maximizing congestion games and valid-utility games [24] that has been studied for the content distribution in service provider networks [14]. Mirrokni and Vetta [18] show that after one round of best responses in which each player get exactly one chance to play best response, players reach an $O(\log n)$-approximate solution.

2 Preliminaries

2.1 General Definitions

Strategic games. A strategic game (or a normal-form game) $\Lambda =< N, (\Sigma_i), (u_i) >$ has a finite set $N = \{1, \ldots, n\}$ of players. Player $i \in N$ has a set Σ_i of actions (or strategies). We call a game *symmetric* if all players share the same set of strategies, otherwise we call it *asymmetric*. The joint action set is $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ and a joint action $S \in \Sigma$ is also called a *profile* or *strategy profile*. The payoff function of player i is $u_i : \Sigma \to \mathbb{R}$, which maps the joint action $S \in \Sigma$ to a real number. Let $S = (S_1, \ldots, S_n)$ denote the profile of actions taken by the players, and let $S_{-i} = (S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_n)$ denote the profile of actions taken by all players other than player i. Note that $S = (S_i, S_{-i})$. An *improvement* move S'_i for a player i in a profile S is a move for which $u_i(S_{-i}, S'_i) \geq u_i(S)$. A *best response* move S''_i for a player i in a profile S is an improvement move that has the maximum payoff. In this paper, we consider two types of games: *cost minimizing games* and *profit maximizing games*. In cost minimizing games, each player i wants to minimize the cost $c_i(S) = -u_i(S)$ in strategy profile S. This type of games include congestion games with polynomial latency functions. In profit maximizing games, each player i wants to maximize the profit $p_i(S) = u_i(S)$ in strategy profile S. This type of games include market sharing games and cut games.

Nash equilibria (NE): A joint action $S \in \Sigma$ is a *pure Nash equilibrium* if no player $i \in N$ can benefit from unilaterally deviating from his action to another action, i.e., $\forall i \in N\ \forall S'_i \in$

$\Sigma_i \ : \ u_i(S_{-i}, S'_i) \leq u_i(S)$. We can also define α-Nash equilibria as follows. For $1 > \alpha > 0$, a state S is an α-Nash equilibrium if for every player i, $c_i(S_{-i}, S'_i) \geq (1-\alpha)c_i(S)$ for all $S'_i \in \Sigma_i$.
State graph. Given any game Λ, the state graph $G(\Lambda)$ is an arc-labelled directed graph as follows. Each vertex in the graph represents a joint action S. There is an arc from state S to state S' with label i iff there exists player i and action $S'_i \in \Sigma_i$ such that $S' = (S_{-i}, S'_i)$, i.e., S' is obtained from S by a move of a single player i that improves his payoff from S to S'.
Exact potential games. A game is called an *exact potential game* if there is a function ϕ such that for any edge of the state graph (S, S') with deviation of player i, we have $\phi(S') - \phi(S) = u_i(S') - u_i(S)$. We denote the minimal potential of the game by ϕ^*.
Social function. Given any game Λ, in order to measure the performance of strategy profiles of players, we define a social function for any strategy profile S. This social function for minimizing cost games is denoted by $cost(S)$ and we denote by $OPT(\Lambda)$ the minimal social cost of a game Λ. i.e., $OPT(\Lambda) = \min_{S\in\Sigma} cost_\Lambda(S)$. We denote by $cost_Z(S)$, the sum of the payoffs of the players in the set Z, when the game Λ is clear from the context, i.e., $cost_Z(S) = \sum_{i\in Z} c_i(S)$. For profit maximizing games, the social function is denoted by $profit(S)$ and we denote by $OPT(\Lambda)$ the maximal social cost of a game Λ. i.e., $OPT(\Lambda) = \max_{S\in\Sigma} profit_\Lambda(S)$. We denote by $profit_Z(S)$, the sum of the payoffs of the players in the set Z, when the game Λ is clear from the context, i.e., $profit_Z(S) = \sum_{i\in Z} p_i(S)$.
α-Nash dynamics. For $0 < \alpha \leq 1$, this dynamics allows only *α-moves* of the players, where α-move of a player is a move that improves his cost by a factor more than $1-\alpha$, i.e., if player i moves from action S_i to action S'_i then $c_i(S_{-i}, S'_i) < (1-\alpha)c_i(S)$. We consider the *unrestricted* α-Nash dynamics with *liveness* property, which allows an adversary to order the players moves in each round as long as every player has at least one chance to move in each round. The liveness property requirement is that in each interval of length T every player appears at least once. For profit maximizing games, an α-move is a move that increases the payoff by a factor more than $1+\alpha$. In these games, we study α-Nash dynamics under the assumption that players play a best response when they get a chance. We call this dynamics, the α-Nash best-response dynamics. Also, an *α-Nash best-response move* is a best response α-move.

α-Nash best-response dynamics is also considered by [17, 8] (called $1+\alpha$-greedy players). The liveness property have been considered by [7] and [18]. Mirrokni and Vetta [18] call a round in which each player gets at least a chance to move, a *covering walk*.
Nice Potential Games. Consider a potential game Λ. Let S be a profile of the players and let S'_i be the best response for any player i. For each player i, let $\Delta_i(S) = c_i(S) - c_i(S_{-i}, S'_i)$ and let $\Delta(S) = \sum_i \Delta_i(S)$. Also, for any set of players Z, let $\Delta_Z(S) = \sum_{i\in Z} \Delta_i(S)$. We may drop the (S) part of the terms and denote these terms by Δ_i and Δ_Z, if the profile is determined clearly in the context.

Definition 2.1 *An exact potential game Λ with potential function ϕ is β-*nice *iff for any state S, it holds that (i) $cost(S) \leq \beta OPT(\Lambda) + 2\Delta(S)$, and (ii) $\phi(S) \leq cost(S)$.*

We consider exact potential games, which are β-nice, where β is the price of anarchy of the game. We show that the α-Nash dynamics converges in polynomial time to a state S with $\Delta(S)$ that is arbitrarily close to zero. Therefore the approximation ratio of the solution S is arbitrarily close to the price of anarchy.
Bounded Jump Property.

Definition 2.2 *(γ-Bounded Jump). For any value $\gamma \geq 1$, a game Λ satisfies the γ-bounded jump condition if for every profile S and every player i with improvement move S'_i, it holds that*

1. $c_j(S) - c_j(S_{-i}, S'_i) \leq c_i(S)$.
2. *for every improvement action S'_j of player j, it holds $c_j(S_{-\{i,j\}}, S'_i, S'_j) - c_j(S_{-j}, S'_j) \leq \gamma \cdot c_i(S_{-i}, S'_i)$.*

Lemma 4.6 shows that congestion games with resources that satisfy the γ-bounded jump condition, studied in [7, 23], satisfy the γ-bounded jump property according to definition 2.2. Therefore it is sufficient to assume the bounded jump property according to definition 2.2 for this class of games.
ε-approximate α-equilibria. Given a strategy profile S, we call the set of players that cannot make an α-move, *α-equilibrium* players.

Definition 2.3 *A state S is an ε-approximate α-equilibrium if $\Delta_O(S) \leq \varepsilon \cdot cost(S)$ where O is the set of players that can play an α-move.*

2.2 Cost Minimizing Congestion Games

In this part, we define cost minimizing congestion games. Since the focus of this paper is on these games, and for brevity, we call these games, congestion games.
Unweighted Congestion Games. An unweighted congestion game is defined by a tuple $< N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E} >$ where E is a set of facilities, $\Sigma_i \subseteq 2^E$ the strategy space of player i, and $d_e : \mathbb{N} \to \mathbb{Z}$ a delay function associated with resource e. For a joint action S, we define the *congestion* $n_e(S)$ on resource e by $n_e(S) = |\{i | e \in S_i\}|$, that is $n_e(S)$ is the number of players that selected an action containing resource e in S. The cost $c_i(S)$ of player i in a joint action S is $c_i(S) = -u_i(S) = \sum_{e \in S_i} d_e(n_e(S))$. [20] showed that every congestion game possesses at least one pure Nash equilibrium by considering the potential function $\phi(S) = \sum_e \sum_{i=1}^{n_e(S)} d_e(i)$.
Weighted Congestion Games. In *weighted congestion games*, player i has weighted demand w_i. We denote by $l_e(S)$, the congestion(load) on resource e in a state S, i.e., $l_e(S) = \sum_{i | e \in S_i} w_i$. The cost of a player in a state S is $c'_i(S) = \sum_{e \in S_i} d_e(l_e(S))$. The total cost is the weighted sum $cost(S) = \sum_{i \in N} w_i c'_i(S) = \sum_{e \in E} l_e d_e(l_e(S))$. Note that congestion games is a special case of weighted congestion games with $w_i = 1$ for every player i. [13] showed that every weighted congestion game with linear latency functions possesses at least one pure Nash equilibrium by considering a potential function equivalent to $\phi(S) = \frac{1}{2}\left(\sum_e l_e(S) d_e(l_e(S)) + \sum_i \sum_{e \in S_i} w_i d_e(w_i)\right)$. We use the fact that this potential function is an exact potential function if the cost of a player in a state S is $w_i c'_i(S)$. To simplify the presentation of the results we assume that the cost of any player i in a state S is $c_i(S) = w_i c'_i(S)$.

2.3 Profit Maximizing Congestion Games

Cut Games. Cut game is a profit maximizing congestion game that is defined on an edge-weighted undirected graph $G(V, E)$, with n vertices and edge weights $w : E(G) \to \mathbb{Q}^+$. We 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 $\mathsf{Adj}(v)$ be the set of neighbors of v. Let also $w_v = \sum_{u \in \mathsf{Adj}(v)} w_{uv}$. 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 $profit(S) = \sum_{v \in T, u \in \bar{T}} w_{uv}$.

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 choose one side of the cut, i.e. v can choose $S_v = -1$ or $S_v = 1$. A strategy profile $S = (S_1, S_2, \ldots, S_n)$, 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 $p_v(S)$, is equal to the contribution of v in the cut, i.e. $p_v(S) = \sum_{i : S_i \neq S_v} w_{iv}$. It follows that the cut value is equal to $\frac{1}{2} \sum_{v \in V} p_v(S)$. If S is clear from the context, we use p_v instead of $p_v(S)$ to denote the payoff of v. We denote the maximum value of a cut in G, by $c(G)$. These games are exact potential games, and the potential function is $\phi(S) = profit(S) = \sum_{v \in T, u \in \bar{T}} w_{uv}$.

Market Sharing Games. A market sharing game is defined by a tuple $< N, M, (\Sigma_i)_{i \in N}, (v_j)_{j \in M} >$ where M is a set of markets, $\Sigma_i \subseteq 2^M$ the strategy space of player i, and v_j the value of market j. For a joint action S, we define the *congestion* $n_j(S)$ on market j by $n_j(S) = |\{i | j \in S_i\}|$, that is $n_j(S)$ is the number of players that selected an action containing market j in S. The payoff $p_i(S)$ of player i in a joint action S is $p_i(S) = u_i(S) = \sum_{j \in S_i} \frac{v_j}{n_j(S)}$. Market sharing games are maximization congestion games with potential function $\phi(S) = \frac{1}{\log n} \sum_{j \in M} \sum_{i=1}^{n_j(S)} \frac{v_j}{i}$. The social function is the sum of payoff of players or the total value of the market satisfied, i.e., $profit(S) = \sum_{i \in N} p_i(S) = \sum_{j \in \cup_{i \in N} S_i} v_j$.

3 Convergence of the α-Nash Dynamics

In this section, we consider the unrestricted α-Nash dynamics with a liveness property for nice exact potential games satisfying the bounded jump property. Throughout this section, let C be the set of α-equilibrium players and let O be the set of all other players, i.e., the players that can make an α-move. First we observe the following simple lemma.

Lemma 3.1 *If a state S is in an ε-approximate α-equilibrium, then $\Delta(S) \leq (\alpha + \varepsilon) cost(S)$.*

Proof: Since C is the set of α-equilibrium players, $\Delta_C(S) \leq \alpha \cdot cost_C(S)$. Thus, $\Delta(S) = \Delta_C(S) + \Delta_O(S) \leq (\alpha + \varepsilon) cost(S)$. ■

As a warmup example, we consider a (restricted) *basic dynamics*, where in each step, among all players that can play an α-move, we choose the player with the maximum absolute improvement, and let him move.

Lemma 3.2 *Let $\frac{1}{8} > \delta \geq \alpha$. Consider an exact potential game Λ that satisfies the* nice *property and any initial state S_{init}. The basic dynamics generates a profile S with $cost(S) \leq \beta(1 + O(\delta)) OPT(\Lambda)$ in at most $O\left(\frac{n}{\delta} \log(\frac{\phi(S_{init})}{\phi^*})\right)$ steps.*

Proof: Consider a step that starts with profile S. Let $\varepsilon_O = \Delta_O(S)/cost(S)$. By definition 2.3 the state S is an ε_O-approximate α-equilibrium. Now, there are two cases:
Case 1: $\varepsilon_O \leq \delta$. It follows from Lemma 3.1 that $\Delta(S) \leq (\alpha + \varepsilon_O) cost(S) \leq (\alpha + \delta) cost(S)$. Hence, by definition 2.1, the dynamics reached $\beta(1 + 4(\alpha + \delta))$-approximation of the optimal cost.

Case 2: $\varepsilon_O > \delta$. It follows that $\Delta_O(S) > \delta \cdot cost(S)$. Hence, there exists a player $j \in O$ such that $\Delta_j(S) > \frac{\delta}{n} cost(S)$. Thus, $\Delta_j(S) > \frac{\delta}{n}\phi(S)$, since $\phi(S) \leq cost(S)$. Therefore the potential gain is at least $\frac{\delta}{n}\phi(S)$. Let $\phi(t)$ denote the potential in step t. Then, $\phi(t) \leq \phi(S_{init})(1-\frac{\delta}{n})^t$. Since $\phi(t) \geq \phi^*$, the upper bound on the number of steps follows. ∎

The above basic Nash dynamics requires some coordination that chooses the player with the maximum gain at each step. Now we show similar results for unrestricted Nash dynamics.

Theorem 3.3 *Let $\frac{1}{8} > \delta \geq 4\alpha$. Consider an exact potential game Λ that satisfies the* nice *property and the γ-bounded jump condition. For any initial state S_{init}, the unrestricted α-Nash dynamics generates a profile S with $cost(S) \leq \beta(1+O(\delta))OPT(\Lambda)$ in at most $O\left(\frac{\gamma n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*})\cdot T\right)$ steps.*

Before proving Theorem 3.3 we point out that the α-Nash dynamics is necessary for polynomial time convergence to nearly optimal solutions for nice exact potential games satisfying the bounded jump property, that is, we show that even after exponentially many steps, the unrestricted exact Nash dynamics with a liveness property for asymmetric congestion games with linear delay functions may generate strategy profiles whose social cost is far from the optimal solution.

Theorem 3.4 *There exists an exact potential game Λ that satisfies the* nice *property and the γ-bounded jump condition, and an initial state S_{init} from which the unrestricted exact best-response dynamics generates a profile S with $cost(S) \geq \Omega(\frac{\sqrt{n}}{\log n})OPT$ after an exponential number of steps. In particular, this holds for a congestion game with linear latency functions.*

The proof of this theorem is based on constructing a long involved example with several components, and is left to Section A of the appendix. We now present the proof of Theorem

Proof: (of Theorem 3.3) Let $\alpha' = 4\alpha$. It is sufficient to consider the case that the players are not in a δ-approximate α'-equilibrium, since otherwise it follows from Lemma 3.1 and Definition 2.1 that the dynamics reached a $\beta(1+4(\alpha'+\delta))$-approximation of the optimal cost. We show that in each interval of T steps the potential decreases by a factor of at least $\frac{\alpha\delta}{4\gamma n}$. Let $S^0, S^1, \ldots, S^T$ denote the joint actions of the players in times $0, 1, \ldots, T$ of this interval respectively. Since S^0 is not a δ-approximate α'-equilibrium, there exists a player with an improvement α'-move. Consider player j with the maximum absolute improvement α'-move and let S'_j be his best response. Recall that $\Delta_j(S^0) = c_j(S^0) - c_j(S^0_{-j}, S'_j)$. Let $\Delta'_j = \Delta_j(S^0)$ and let t' be the first time in this interval that player j is allowed to move. We denote by U the set of times before time t', where players made α-moves and we denote by $w(t)$ the player that moved at time t for each $t \in U$. Let $A = \sum_{t\in U} c_{w(t)}(S^t)$ be the sum of the costs of the moving players when they make their moves. Now, we consider two cases:
Case 1: $A \leq \frac{\Delta'_j}{4\gamma}$. By the first condition of the bounded jump property, we have for each $t \in U$

$$c_j(S^t) - c_j(S^{t+1}) \leq c_{w(t)}(S^t). \tag{1}$$

Summing over all times $t \in U$, we obtain:

$$c_j(S^0) - c_j(S^{t'}) \leq \sum_{t\in U} c_{w(t)}(S^t) = A \leq \frac{\Delta'_j}{4\gamma} \leq \frac{\Delta'_j}{4}. \tag{2}$$

Where the first inequality follows since the sum of the left hand side of equation (1) telescopes.

Similarly, by the second property of the bounded jump assumption, we obtain

$$c_j(S^{t'}_{-j}, S'_j) - c_j(S^0_{-j}, S'_j) \leq \gamma \cdot A \leq \gamma \frac{\Delta'_j}{4\gamma} \leq \frac{\Delta'_j}{4}. \tag{3}$$

By summing inequalities (2) and (3), we get

$$c_j(S^{t'}) - c_j(S^{t'}_{-j}, S'_j) \geq c_j(S^0) - c_j(S^0_{-j}, S'_j) - \frac{\Delta'_j}{2} = \Delta'_j - \frac{\Delta'_j}{2} = \frac{\Delta'_j}{2}. \tag{4}$$

By the second property of the bounded jump assumption we also get

$$c_j(S^{t'}) \leq c_j(S^0) + \gamma \cdot A \leq c_j(S^0) + \gamma \frac{\Delta'_j}{4\gamma} = c_j(S^0) + \frac{\Delta'_j}{4}. \tag{5}$$

Hence,

$$c_j(S^{t'}) \leq c_j(S^0) + \frac{\Delta'_j}{4} < \frac{\Delta'_j}{\alpha'} + \frac{\Delta'_j}{4} < 2\frac{\Delta'_j}{4\alpha} = \frac{\Delta'_j}{2\alpha}.$$

Where the second inequality follows from the fact that Δ'_j is the improvement of player j when making his best response, which is an α'-move in step 0. Thus, $\alpha \cdot c_j(S^{t'}) < \frac{\Delta'_j}{2}$. As a result, using this inequality and inequality (4), we get $\alpha \cdot c_j(S^{t'}) < c_j(S^{t'}) - c_j(S^{t'}_{-j}, S'_j)$. Therefore, player j can make an α-move at time t' and decrease the potential ϕ by at least $\alpha \cdot c_j(S^{t'}) \geq \alpha \frac{\Delta'_j}{2} \geq \frac{\alpha\delta}{2n}\phi(S^0)$.

Case 2: $A > \frac{\Delta'_j}{4\gamma}$. Since A is the sum of the costs of players making an α-move when making the move, these players decrease the potential ϕ by at least $\alpha A > \frac{\alpha \Delta'_j}{4\gamma} \geq \frac{\alpha\delta}{4\gamma n}\phi(S^0)$.

Let $\phi(i)$ denote the potential in round i. Then, in both cases $\phi(i) \leq \phi(S_{init})(1 - \frac{\alpha\delta}{4\gamma n})^i$. Since $\phi(i) \geq \phi^*$, the upper bound on the number of steps follows. ■

Remark 3.5 The above theorem shows that we reach a state with cost at most $\beta(1 + O(\delta))$ of the optimum after polynomial number of α-moves. Eventhough after this state the cost of solutions can increase, it follows from the proof of the theorem that the number of states in which the cost of the solution is more than a $\beta(1 + O(\delta))$-approximation is at most $O(\frac{\gamma n}{\alpha\delta} \log(\frac{\phi(S_{init})}{\phi^*})T)$. In addition, since the potential function is always decreasing after any α-move, the cost can increase by a factor of at most $\frac{cost(S)}{\phi(S)}$. It is not hard to show that the ratio $\frac{cost(S)}{\phi(S)}$ for any strategy profile in congestion games with polynomial delay functions of degree d is at most $O(d)$ and for weighted congestion games with linear functions is at most $O(1)$. As a result, for both type of congestion games that we consider in Section 4, the cost of any state after a polynomial number of steps reach a constant-factor approximate solution and remains within a constant factor of the optimal solution.

4 Congestion Games

In this section we consider weighted congestion games with linear latency functions and congestion games with linear and polynomial latency functions.

4.1 Linear Latency Functions

In this section we consider weighed and unweighted congestion games with linear latency functions. Specifically $d_e(x) = a_e x + b_e$ for each resources $e \in E$, where a_e and b_e are nonnegative reals. For simplicity we only consider the identity function $d_e(x) = x$. It is easy to verify that all the proofs work for the general case as well.

4.1.1 Weighted Congestion Games

We first show that weighted congestion games with linear latency functions are β-*nice* according to definition 2.1 with $\beta = \frac{3+\sqrt{5}}{2} \approx 2.618$.

Lemma 4.1 *Congestion games with linear latency functions are* β-nice *potential games with* $\beta = \frac{3+\sqrt{5}}{2}$.

Next we show that weighted congestion game with linear delay functions satisfy the 1-bounded jump condition.

Lemma 4.2 *Let* Λ *be a weighted congestion game with linear delay functions. Then, the game* Λ *satisfies the* 1*-bounded jump condition according to definition 2.2.*

The proof of the above two lemmas can be found in the appendix. Theorem 3.3 and Lemmas 4.1, 4.2 yield the following corollary.

Corollary 4.3 *Let* $\frac{1}{8} > \delta \geq \alpha$. *Consider a weighted congestion game* Λ *with linear latency functions and any initial state* S_{init}. *The unrestricted* α*-Nash dynamics with liveness property generates a profile* S *with* $cost(S) \leq \frac{3+\sqrt{5}}{2}(1+O(\delta))OPT(\Lambda)$ *in at most* $O\left(\frac{n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*})\cdot T\right)$ *steps.*

4.1.2 Unweighted Congestion Games

We first show that congestion games with linear latency functions are β-*nice* according to definition 2.1 with $\beta = 2.5$. In the proof of this lemma, we use two Lemmas which appear in [9], and are stated in the appendix.

Lemma 4.4 *Congestion games with linear latency functions are* β-nice *potential games with* $\beta = 2.5$.

Next we show that unweighted congestion games with resources that satisfy the γ-bounded jump condition, satisfy the γ-bounded jump condition according to definition 2.2.

Definition 4.5 *(resource* γ*-bounded jump). Resource* e *satisfies the* γ*-bounded jump condition if its delay function satisfies* $d_e(x+1) \leq \gamma \cdot d_e(x)$ *for every* $x \geq 1$*, for* $\gamma \geq 1$.

Lemma 4.6 *Let* Λ *be a congestion game with nonnegative, non-decreasing delay functions in which every resource has* γ*-bounded jump. Then, the game* Λ *satisfies the* γ*-bounded jump condition according to definition 2.2.*

The proof of the above two lemmas can be found in the appendix. Theorem 3.3, Lemmas 4.4, 4.6 and the fact that resource with linear latency function has 2-bounded jump, yield the following corollary.

Corollary 4.7 *Let $\frac{1}{8} > \delta \geq \alpha$. Consider a congestion game Λ with linear latency functions and any initial state S_{init}. The unrestricted α-Nash dynamics with liveness property generates a profile S with cost$(S) \leq 2.5(1 + O(\delta))OPT(\Lambda)$ in at most$O\left(\frac{n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*})\cdot T\right)$ steps.*

4.2 Polynomial Latency Functions

In this section, we consider congestion games with polynomial latency functions of degree d. We show that congestion games with polynomial latency functions are β-*nice* according to definition 2.1 with $\beta = d^{d(1-o(1))}$. Price of anarchy results which appear in [9] imply that for $\beta = d^{d(1-o(1))}$ and for every profile S equation (**??**) in definition 2.1 holds.

Lemma 4.8 *Congestion games with polynomial latency functions of degree d are β-*nice *potential games with $\beta = d^{d(1-o(1))}$.*

Theorem 3.3, Lemmas 4.8, 4.6 and the fact that resource with polynomial of degree d latency function has 2^d-bounded jump, yield the following corollary.

Corollary 4.9 *Let $\frac{1}{8} > \delta \geq \alpha$. Consider a congestion game Λ with polynomial latency functions of degree d and any initial state S_{init}. The unrestricted dynamics generates a profile S with cost$(S) \leq d^{d(1-o(1))}(1 + O(\delta))OPT(\Lambda)$ in at most $O\left(\frac{2^d \cdot n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*})\cdot T\right)$ steps.*

5 Profit Maximizing Congestion Games

In this section, we extend the results for cost minimizing congestion games to profit maximizing congestion games. We first define some preliminaries for these games. Consider an exact potential game Λ. Let S be a profile of the players and let $S_i^{'}$ be a best response strategy for player i in strategy profile S. The payoff of player i in strategy profile S is denoted by $p_i(S)$ and each player wants to maximize its payoff. In this setting, for each player i, let $\Delta_i(S) = p_i(S_{-i}, S_i') - p_i(S)$ and let $\Delta(S) = \sum_i \Delta_i(S)$.

Definition 5.1 *An exact potential game Λ with potential function ϕ is β-*nice *iff for any state S it holds that (i) $\beta \cdot (profit(S) + \Delta(S)) \geq OPT(\Lambda)$, and (ii) $\phi(S) \leq profit(S)$.*

Definition 5.2 *(γ-Bounded Jump). Consider any profile S and any player i with improvement move S_i'. Then, for every player j the following properties hold:*

1. $p_j(S_{-i}, S_i') - p_j(S) \leq p_i(S_{-i}, S_i')$
2. *for every improvement action S_j' of player j, it holds* $p_j(S_{-j}, S_j') - p_j(S_{-\{i,j\}}, S_i', S_j') \leq \gamma \cdot p_i(S_{-i}, S_i')$

5.1 Convergence of profit maximizing games

Similar to the proof of Theorem 3.3 for convergence of unrestricted α-Nash dynamics in cost minimizing games, we can prove the following general theorem for convergence time of the α-Nash best-response dynamics profit maximizing games.

Theorem 5.3 *Let $\frac{1}{8} > \delta \geq 4\alpha$. Consider an exact potential game Λ that satisfies the* nice *property and the bounded jump condition. For any initial state S_{init} the unrestricted α–Nash best-response dynamics with liveness property generates a profile S with $\beta(1+O(\delta))profit(S) \geq OPT(\Lambda)$ in at most $O\left(\frac{\gamma n}{\alpha\delta}\log(\frac{\phi^*}{\phi(S_{init})})\cdot T\right)$ steps.*

The proof of this theorem is very similar to that of Theorem 3.3 and is left to Section C in the appendix.

6 Cut Games and Market Sharing Games

Using Theorem 5.3, in order to prove polynomial convergence of α-Nash best-response dynamics in cut games and market sharing games, we can show that both of these games satisfy the 2-nice and 1-bounded jump properties. The proofs of these properties can be found in Sections D and E of the appendix.

Corollary 6.1 *Let $\frac{1}{8} > \delta \geq 4\alpha$. Consider a cut game or a market sharing game Λ with and any initial state S_{init}. The unrestricted α–Nash best-response dynamics with liveness property generates a profile S with profit $\frac{1}{(2+O(\delta))}OPT(\Lambda)$ in at most $O\left(\frac{n}{\alpha\delta}\log(\frac{\phi^*}{\phi(S_{init})})\cdot T\right)$ steps.*

References

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 613–622, 2006.

[2] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304, 2004.

[3] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proc. 37th ACM Symp. on Theory of Computing*, pages 57–66, 2005.

[4] M. Charikar, C. Mattieu, H. Karloff, S. Naor, and M. Saks. Best response dynamics in multicast cost sharing. 2007. Unpublished Manuscript.

[5] C. Chekuri, J. Chuzhoy, L. Lewin, S. Naor, and A. Orda. Non-cooperative multicast and facility location games. In *ACM EC*, pages 72–81, 2006.

[6] Xi Chen and Xiaotie Deng. Settling the complexity of two-player nash equilibrium. In *Proc. 47nd IEEE Symp. on Found. of Comp. Science*, pages 261–272, 2006.

[7] Steve Chien and Alistair Sinclair. Convergence to approximate nash equilibria in congestion games. In *Proc. 18th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2007.

[8] G. Christodolou, V. S. Mirrokni, and A. Sidiropolous. Convergence and approximation in potential games. In *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 2006.

[9] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *37th ACM Symposium on Theory of Computing*, pages 67–73, Baltimore, MD, USA, 22–24 May 2005.

[10] Artur Czumaj and Berthold Vöcking. Tight bounds for worst-case equilibria. In *SODA*, pages 413–420, 2002.

[11] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–78, 2006.

[12] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 604–612, 2004.

[13] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science, Special Issue on ICALP 2004*, 348:226–239, 2005.

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

[15] M. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS*, 2005.

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

[17] V. S. Mirrokni. *Approximation Algorithms for Distributed and Selfish Agents*. Massachusetts Institute of Technology, 2005.

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

[19] Christos Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of 33rd STOC*, pages 749–753, 2001.

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

[21] T. Roughgarden and Eva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259, 2002.

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

[23] A. Skopalik and B. Vöcking. Inapproximability of convergence in congestion games. 2007. Unpublished Manuscrip.

[24] 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.

A Convergence of the exact Nash Dynamics

Here, we consider the unrestricted exact Nash dynamics with a liveness property for nice exact potential games satisfying the bounded jump property. We show that even after exponentially many steps, the exact Nash dynamics for linear congestion games may generate strategy profiles whose social cost is far from the optimal solution.

Theorem A.1 *After exponentially many steps of the exact Nash dynamics of congestion games with linear latency functions, we may reach solutions whose cost is more than* $\Omega(\frac{\sqrt{n}}{\log n})OPT$.

Proof: We construct a congestion game Γ_n consisting of four types of players: (i) *counter players* representing a binary counter, (ii) *congestion players* consisting of a group of A-players and one C-player, (iii) *bit players* denoted by B-players, and (iv) *trigger players* consisting of T-players, R-players, and a Q-player. For a large integer number M, we show that there exists a best-response sequence of exponential length in which the the counter players have delays of at least M. However, each of the counter players has a strategy with one unique *extra* resource with the delay function $\ell(x) = \frac{5M}{\sqrt{n}}x$. But each time a counter player has the chance to change his strategy, his extra resource is congested by $\sqrt{n}$ *congestion* players. In each step of the counter, the congestion players successively allocate all extra resources of the counter players. Thus, after each step, every counter player gets a chance to deviate but it does not have an incentive to change to its extra resource. However, in the optimal solution, the counter players can deviate to their extra strategy that has delay of $\frac{M}{\sqrt{n}}$. The cost of the optimal solution is dominated by the cost of the $\sqrt{n}$ congestion players that have delay of $O(\log n)M$. Thus, even after an exponentially long sequence of best response, the social cost is at least $\Omega(\frac{n}{\sqrt{n}\log n})OPT = \Omega(\frac{\sqrt{n}}{\log n})OPT$.

First we describe the high-level idea of the construction of the congestion game. A main component of this game is an *n-bit counter* consisting of $4n$ counter players. This counter is similar to existing examples, e.g., [**?**]. The best response sequence of the counter players count downwards from $2^n - 1$ to 0. In each counting step of the counter, we give all players a chance to move. For each player, there exists an extra a-resource and an additional Alt strategy that consists of only this resource. If a player changes to that strategy and no other player is on that resource, he can decrease his delay by a factor of $k = \sqrt{n}$. However, there are $k = \sqrt{n}$ congestion players, denoted by A-players, that are on that resource making a deviation of a counter player not favorable In addition, in each counting step of the counter, we let the C-player successively occupy the c-resources of the bits of the counter that are 1. This prevents these bits from switching to 0. Thus, we can give all counter players a chance

Player	Strategies	Resources	Delays
Init_i	One	r_i^1	$4M + 8\delta^i$
	Zero	r_i^2	$M/M + 2\delta^i$
		r_i^3	$M/M + 4\delta^i$
		c_i	$M/2M$
		r_i^5	$M/M + 9\delta^i$
	Alt	a_i^1	$\frac{5M}{k}x$
Change_i	One	r_i^2	$M/M + 2\delta^i$
		r	M
	Zero	r_i^4	$M/M + 3\delta^i$
		q	$100n(n+1)x$
		r_{i-1}^6	$M + 10\delta^{i-1}$
	Alt	a_i^2	$\frac{5M}{k}x$

Player	Strategies	Resources	Delays
Done_i	One	r_i^3	$M/M + 4\delta^i$
	Zero	r_i^4	$M/M + 3\delta^i$
	Alt	a_i^3	$\frac{5M}{k}x$
Reset_i	One	r_i^6	$M + 10\delta^i$
		r	M
	Zero	r_i^5	$M + 9\delta^i$
		r_{i-1}^6	$M + 10\delta^{i-1}$
	Alt	a_i^4	$\frac{5M}{k}x$

Figure 1: The strategies of the four counter players for the i-th bit. For delay functions, if they have x in the description, the delay function $f(x)$ is given in terms of the congestion x. If the resource is used only by one or two players, the delay function is denoted by r_1/r_2 where r_1 and r_2 are delays for congestion 1 and 2 respectively. The scaling factor δ is at least $200n(n+1)$.

to move. To ensure that all the A-players pile up on the same a-resources, we have $log 4n$ B-players that encode binary numbers corresponding to the a-resources which are to be allocated by the A-players. This way, we make sure that the only profitable deviation of the A-players is to allocate this particular a-resource, since the delay of the other strategies is higher due to the B-players. Finally, the delay of the trigger players increases by the counter players in each step of the counter. Their best responses increase the delays of the bit and congestion players such that the all the aforementioned strategy changes are best responses.

We now describe the construction of the game. We say a player is activated or we activate a player if we let him play his best response. We first describe the details of the n-bit counter. For each bit i with $1 \leq i \leq n$ we have 4 players, see Figure A for the complete description of their strategies and of the resources.

We say the i-th bit of the counter is 1 if player Init_i plays his One-strategy and 0 otherwise. We construct a best-response sequence that consist of exponentially many rounds. We start the sequence with all counter players playing One and we ensure that no counter player changes to Alt until the end of the process. In each round, all players are activated at least once.

In order to prove the result, we make certain assumptions on the usage of certain resources by other players. We will later show, that these assumptions hold throughout the process. Let x be the value of the counter and i be the bit that flips from 1 to 0 when changing to $x - 1$. Throughout the process, we make sure that the following three conditions hold:

1. (A1) Each time we activate a player, the a-resource in his Alt strategy is congested by k A-players.
2. (A2) If we activate a player $\text{Init}_{i'}$ with $i' \neq i$ and $\text{Init}_{i'}$ plays One, then the resource $b_{i'}$ is allocated by another player.

3. (A3) The resource q is allocated by at most one other player.

	activated player	best response
1	$\text{Init}_{i+1},\ldots,\text{Init}_n$	unchanged
2	$\text{Done}_{i+1},\ldots,\text{Done}_n$	One if his Init player plays One, Zero otherwise.
3	$\text{Change}_{i+1},\ldots,\text{Change}_n$	One
4	$\text{Reset}_i,\ldots,\text{Reset}_n$	One
5	Init_i	Zero
6	Trigger_i	Zero
7	$\text{Reset}_{i-1},\ldots,\text{Reset}_1$	Zero
8	$\text{Init}_{i-1},\ldots,\text{Init}_1$	One
9	$\text{Done}_{i-1},\ldots,\text{Done}_1$	One
10	$\text{Trigger}_{i-1},\ldots,\text{Trigger}_1$	One
11	Done_i	Zero
12	Trigger_i	One

Figure 2: Sequence of activations and best responses of one round four counter players. This round corresponds to one step of the counter in which the i-th bit switches from 1 to 0.

The activation sequence of counter players is described in Figure A. Essentially, the players corresponding to bits greater than i do not change their strategies. The player Init_i changes to Zero. This results in a sequence of best responses in which the less significant bits change to 1. Thus, under the above assumptions, the value of the counter decreases by exactly one in each step. Note, that during each step one trigger player allocates the resource T and leaves it again. We will now describe the next component, a set of trigger players that make use of this fact.

The trigger component consists of a large set of trigger players. We use these player to repeatedly increase and decrease the delay on some resources of the components that we describe later. Trigger players consist of one player Q and several T-players and R-players. Each trigger player has two strategies Wait and Trigger. See Figure A for a detailed description of the strategies and resources. Only player Q is interested in resource q. Hence, condition A3 is satisfied.

The best response for Q is Wait, if none of the players of the counter allocates the resource q, otherwise his best response is Trigger. The best response for a R-player is Wait, if Q is on Wait, otherwise his best response is Trigger. The best response for a T-player is Wait if Q is on Wait or the R-player with the same index is on Trigger, otherwise his best response is Trigger.

Assume all players play Wait except player Q who plays Trigger. Then there is a best-response sequence that can be divided to $n+1$ segments. Each segment contains contains strategy profiles $S_1,\ldots,S_6$ with S_i occurs before S_{i+1}. These strategy profiles have the following properties:

(S_1) For an arbitrary set $I \subseteq \{0,\ldots,n\}$, each resource in $\{t_{C_i}\}_{i\in I}$ is allocated by one T-player. Each resource in $\{t_{C_i}\}_{i\notin I}$ is not allocated by any T-player.

(S_2) No resource t_{A_j} is allocated by any T-player.

Player	Strategies	Resources	Delays
Q	Wait	q	$100n(n+1)x$
	Trigger	r_T	$100n(n+1)$
		$q^j_{C_i}$	$5x$
		$p^j_{C_i}$	$7x$
		q_{A_i}	$5kx$
		p_{A_i}	$7kx$
		$p_{B^0_i}$	$7mx$
		$p_{B^1_i}$	$7mx$
		$q_{B^0_i}$	$5mx$
		$q_{B^1_i}$	$5mx$
$T^j_{C_i}$ for each $1 \leq i \leq n$ and $1 \leq j \leq n+1$	Wait	$q^j_{C_i}$	$5x$
	Trigger	t_{C_i}	$2x$
		$r^j_{C_i}$	$5x$
$R^j_{C_i}$ for each $1 \leq i \leq n$ and $1 \leq j \leq n+1$	Wait	$p^j_{C_i}$	$7x$
	Reset	$r^j_{C_i}$	$5x$
			3
T_{A_i} for each $1 \leq i \leq n+1$	Wait	q_{A_i}	$5kx$
	Trigger	t_{A_j} for all $1 \leq j \leq k$	$2x$
		r_{A_i}	$5kx$
R_{A_i} for each $1 \leq i \leq n+1$	Wait	p_{A_i}	$7kx$
	Reset	r_{A_i}	$5kx$
			$3k$
$T_{B^0_i}$ for each $1 \leq i \leq n+1$	Wait	$q_{B^0_i}$	$5mx$
	Trigger	$t_{B^0_j}$ for all $1 \leq j \leq m$	$2x$
		$r_{B^0_i}$	$5mx$
$R_{B^0_i}$ for each $1 \leq i \leq n+1$	Wait	$p_{B^0_i}$	$7mx$
	Reset	$r_{B^0_i}$	$5mx$
			$3k$
$T_{B^1_i}$ for each $1 \leq i \leq n+1$	Wait	$q_{B^1_i}$	$5mx$
	Trigger	$t_{B^1_j}$ for all $1 \leq j \leq m$	$2x$
		$r_{B^1_i}$	$5mx$
$R_{B^1_i}$ for each $1 \leq i \leq n+1$	Wait	$p_{B^1_i}$	$7mx$
	Reset	$r_{B^1_i}$	$5mx$
			$3k$

Figure 3: Description of all trigger players. The description of delay functions $f(x)$ is given in terms of the congestion x.

(S_3) Each resource $t_{B_j^0}$ is allocated by one T-player. No resource $t_{B_j^1}$ is allocated by any T-player.

(S_4) Each resource $t_{B_j^1}$ is allocated by one T-player. No resource $t_{B_j^0}$ is allocated by any T-player.

(S_5) Each resource t_{A_j} is allocated by one T-player.

(S_6) No resource t_{C_i} is allocated by any T-player.

Note that during each step of the counter, the resource q is allocated by a player of the counter. At the end of each step, it is no longer allocated by any player of the counter. We therefore assume that during each step of the counter, there exist a best-response sequence of the trigger players consisting of $n+1$ of the aforementioned segments.

We are now ready to describe the remaining component consisting of bit players and congestion players. In this part, we make use of sequence of strategy profiles of the trigger players and show that the desired assumptions for the counter are met. There are $m = \lceil \log 4n \rceil$ bit players, denoted by B-players, that encode which of the a-resources are to be allocated by $k = \lceil \sqrt{n} \rceil$ A-players. Furthermore, there is one C-player that allocates one of the c-resources. The complete description of these players and their strategies can be found in Figure A.

Player	Strategies	Resources	Delays
A_j	Zero	r_A	$6mM + \frac{5M}{k} j$
for each $1 \le j \le k$		t_{A_j}	$2x$
	(i, l)	a_i^j	$\frac{5M}{k} x$
	for each $1 \le i \le n$	b_y^0 if the y-th bit of $4i + l$ is 0	$6Mx$
	and $1 \le l \le 4$	b_y^1 if the y-th bit of $4i + l$ is 1	$6Mx$
		r	3
B_j	Zero	b_j^0	$3Mx$
for each $1 \le j \le m$		$t_{B_j^0}$	$2x$
	One	b_j^1	$3Mx$
		$t_{B_j^1}$	$2x$
C	Zero	r_C	M
		t_{C_0}	$2x$
	i	c_i	Mx
	for each $1 \le i \le n$	t_{C_i}	$2x$

Figure 4: The A-, B-, and C-players and their strategies. The description of delay functions $f(x)$ is given in terms of the congestion x.

Now we describe another condition involving B-players that is met through the process.

Condition C1 No other player except the B-players allocate any b-resource.

Consider a strategy profile corresponding to S_3. In this case, the best response of every B-player is One. Consider a strategy profile corresponding to S_4. Then the best response of every B-player is Zero. Thus, in each of the $n+1$ segments of the best-response sequence of the Trigger players, there is a best response sequence of the B-players that leads to a strategy profile in which the B-player encode an arbitrary number.

We now show, that given a segment of best-responses by the Trigger players there is a best response sequence such that all k A-players allocate an arbitrary resource a_i^j We start with a strategy profile S_2 and activate the A players in ascending order starting with A_1. Each player's best response is Zero. This satisfies condition C1. In the profile S_3 and S_4 the B-players encode the number $i * 4 + j$. In S_5 we active the A players in ascending order starting with A_1. Each player's best response is (i, j). Thus, in each step of the counter, the A-players successively allocate $n+1$ arbitrary a-resources. We choose the a-resources such that condition A1 is satisfied.

Finally, there is exists C-player that is used to make sure that condition A2 is satisfied. Given a strategy profile S_1 in which exactly one resource t_{C_i} is not allocated by a trigger player, his best response is the strategy i. Thus, there is a best-response sequence in which in every step of the counter the C-player successively allocate $n + 1$ arbitrary c-resources. We choose the latency functions of c-resources in a way to make sure that condition A2 is satisfied.

■

B Missing Proofs of Section 4

Proof of Lemma 4.2 The proof require the following two Lemmas. The first lemma appears in [3] and the second lemma is a simple fact.

Lemma B.1 *Consider a weighted congestion game Λ with linear delay functions. Let S be any profile and S^* be a profile of the optimal solution, then*

$$\sum_i c_i(S_{-i}, S_i^*) \leq \sqrt{cost(S)}\sqrt{cost(S^*)} + cost(S^*).$$

Lemma B.2 *For every pair of nonnegative integers x, y, if $x^2 \leq x+1+y$, then $x^2 \leq \frac{3+\sqrt{5}}{2}+2y$.*

Proof: Let S^* be a profile of the optimal solution and let S be any profile. Applying Lemma B.1, we get $\sum_i c_i(S_{-i}, S_i^*) \leq \sqrt{cost(S)}\sqrt{cost(S^*)} + cost(S^*)$. Note that $cost(S) - \sum_i c_i(S_{-i}, S_i^*) \leq \Delta(S)$, since for any player i with best response S_i', $c_i(S_{-i}, S_i') \leq c_i(S_{-i}, S_i^*)$. Thus, by adding $\Delta(S)$ to both sides of the inequality, we get $cost(S) \leq \sqrt{cost(S)}\sqrt{cost(S^*)} + cost(S^*) + \Delta(S)$. Let $x = \sqrt{\frac{cost(S)}{cost(S^*)}}$ and let $y = \frac{\Delta(S)}{cost(S^*)}$. Now, we divide the above inequality by $cost(S^*)$ and express the result in terms of x and y. Thus, $x^2 \leq x+1+y$. Applying Lemma B.2, we get $x^2 \leq \frac{3+\sqrt{5}}{2} + 2y$. This completes the proof of the Lemma. ■

Proof of Lemma 4.2

Proof: Consider any profile S and any player i with improving action S'_i. We first show the first property in definition 2.2. Consider any player j. Then,

$$\begin{aligned} c_j(S) - c_j(S_{-i}, S'_i) &\leq w_j \sum_{e \in (S_i \setminus S'_i) \cap S_j} l_e(S) - (l_e(S) - w_i) \\ &= w_j \sum_{e \in (S_i \setminus S'_i) \cap S_j} w_i = w_i \sum_{e \in (S_i \setminus S'_i) \cap S_j} w_j \\ &\leq w_i \sum_{e \in (S_i \setminus S'_i) \cap S_j} l_e(S) \leq w_i \sum_{e \in S_i} l_e(S) \\ &= c_i(S). \end{aligned}$$

For the second property in definition 2.2. Consider any player j with action S'_j. Then,

$$\begin{aligned} c_j(S_{-\{i,j\}}, S'_i, S'_j) - c_j(S_{-j}, S'_j) &\leq w_j \sum_{e \in (S'_i \setminus S_i) \cap S'_j} (l_e(S_{-j}, S'_j) + w_i) - l_e(S_{-j}, S'_j) \\ &= w_j \sum_{e \in (S'_i \setminus S_i) \cap S'_j} w_i = w_i \sum_{e \in (S'_i \setminus S_i) \cap S'_j} w_j \\ &\leq w_i \sum_{e \in (S'_i \setminus S_i) \cap S'_j} l_e(S_{-i}, S'_i) \leq w_i \sum_{e \in S'_i} l_e(S_{-i}, S'_i) \\ &= c_i(S_{-i}, S'_i). \end{aligned}$$

■

Proof of Lemma 4.4 To present the proof of Lemma 4.4, we need requires the following two Lemmas which appear in [9].

Lemma B.3 *Consider a congestion game Λ with nonnegative, non-decreasing delay functions. Let S be any profile and let S^* be a profile of the optimal solution, then*

$$\sum_i c_i(S_{-i}, S^*_i) \leq \sum_{e \in E} n_e(S^*) d_e(n_e(S) + 1).$$

Lemma B.4 *For every pair of nonnegative integers x, y, it holds $x(y+1) \leq \frac{5}{3}x^2 + \frac{1}{3}y^2$.*

Proof: Let S^* be a profile of the optimal solution and let S be any profile. Applying Lemma B.3, we obtain $\sum_i c_i(S_{-i}, S^*_i) \leq \sum_{e \in E} n_e(S^*) d_e(n_e(S) + 1)$. Applying Lemma B.4, we get

$$\begin{aligned} \sum_i c_i(S_{-i}, S^*_i) &\leq \sum_{e \in E} \left(\frac{5}{3} n_e(S^*)^2 + \frac{1}{3} n_e(S)^2 \right) = \frac{5}{3} \sum_{e \in E} n_e(S^*)^2 + \frac{1}{3} \sum_{e \in E} n_e(S)^2 \\ &= \frac{5}{3} cost(S^*) + \frac{1}{3} cost(S). \end{aligned}$$

Recall that $cost(S) - \sum_i c_i(S_{-i}, S^*_i) \leq \Delta(S)$, where S'_i is the best response of any player i. Thus, by multiplying the inequality by 3/2, adding $\Delta(S)$ to both sides and rearranging the terms, we get $\sum_i c_i(S_{-i}, S^*_i) \leq 2.5 \cdot cost(S^*) + \frac{\Delta(S)}{2}$. Therefore, $cost(S) \leq 2.5 \cdot cost(S^*) + \frac{3}{2}\Delta(S)$.
■

Proof of Lemma 4.6

Proof: Consider any profile S and any player i with improving action S'_i. We first show the first property in definition 2.2. Consider any player j. Then,

$$c_j(S) - c_j(S_{-i}, S'_i) \leq \sum_{e \in S_i \cap S_j} d_e(n_e(S)) \leq c_i(S).$$

For the second property in definition 2.2. Consider any player j with action S'_j. Then,

$$\begin{aligned} c_j(S_{-\{i,j\}}, S'_i, S'_j) - c_j(S_{-j}, S'_j) &\leq \sum_{e \in (S'_i \setminus S_i) \cap S'_j} d_e(n_e(S_{-i}, S'_i) + 1) \\ &\leq \sum_{e \in (S'_i \setminus S_i) \cap S'_j} \gamma \cdot d_e(n_e(S_{-i}, S'_i)) \\ &\leq \gamma \cdot c_i(S_{-i}, S'_i). \end{aligned}$$

Where the second inequality uses the assumption that each resource e has γ-bounded jump. ∎

C Convergence in Profit Maximizing Games

In this section, we prove general convergence results for profit maximizing games. First, we give some definitions. Throughout this section, let C be the set of players that cannot make an α-move; we call these players α-equilibrium players, and let O be the set of all other players, i.e., the players that can make an α-move. Let $\Delta_C(S) = \sum_{i \in C} \Delta_i(S)$ and let $\Delta_O(S) = \sum_{i \in O} \Delta_i(S)$.

Definition C.1 *A state S is a ε-approximate α-equilibrium if $\Delta_O(S) \leq \varepsilon \cdot profit(S)$.*

Now, we observe the following simple lemma.

Lemma C.2 *If a state S is in ε-approximate α-equilibrium, then $\Delta(S) \leq (\alpha + \varepsilon) profit(S)$.*

Proof: Since C is the set of players in α-equilibrium, $\Delta_C(S) \leq \alpha \cdot profit_C(S)$. Thus,

$$\Delta(S) = \Delta_C(S) + \Delta_O(S) \leq (\alpha + \varepsilon) profit(S).$$

∎

As a warmup example, we prove the following lemma about the (restricted) basic best-response dynamics in which at each step we choose a player that can play a best-response α-move. We denote the maximal potential of the game by ϕ^*.

Lemma C.3 *Let $\frac{1}{8} > \delta \geq \alpha$. Consider an exact potential game Λ that satisfies the* nice *property and any initial state S_{init}. The basic best-response dynamics generates a profile S with $\beta(1 + O(\delta)) profit(S) \geq OPT(\Lambda)$ in at most $O(\frac{n}{\delta} \log(\frac{\phi^*}{\phi(S_{init})})$ steps.*

Proof: Consider a step that starts with profile S. Let $\varepsilon_O = \frac{\Delta_O(S)}{profit(S)}$. By definition C.1, the state S is in ε_O-approximate α-equilibrium. Now, there are two cases:
Case 1: $\varepsilon_O \leq \delta$. It follows from Lemma C.2 that $\Delta(S) \leq (\alpha+\varepsilon_O)profit(S) \leq (\alpha+\delta)profit(S)$. Hence, by definition 5.1, the dynamics reached $\beta(1+\alpha+\delta)$-approximation of the optimal cost.
Case 2: $\varepsilon_O > \delta$. It follows that $\Delta_O(S) > \delta \cdot profit(S)$. Hence, there exists a player $j \in O$ such that $\Delta_j(S) > \frac{\delta}{n} profit(S)$. Thus, $\Delta_j(S) > \frac{\delta}{n}\phi(S)$, since $\phi(S) \leq profit(S)$.

Therefore the potential gain is at least $\frac{\delta}{n}\phi(S)$. Let $\phi(t)$ denote the potential in step t. Then, $\phi(t) \geq \phi(S_{init})(1+\frac{\delta}{n})^t$. Since $\phi(t) \leq \phi^*$, the upper bound on the number of steps follows. ■

The following Theorem is the main result of this section.

Theorem C.4 *Let $\frac{1}{8} > \delta \geq 4\alpha$. Consider an exact potential game Λ that satisfies the* nice *property and the bounded jump condition. For any initial state S_{init} the unrestricted α-Nash best-response dynamics with liveness property generates a profile S with $\beta(1+O(\delta))profit(S) \geq OPT(\Lambda)$ in at most $O\left(\frac{\gamma n}{\alpha\delta}\log(\frac{\phi^*}{\phi(S_{init})})\cdot T\right)$ steps.*

Proof: Let $\alpha' = 4\alpha$. It is sufficient to consider the case that the players are not in a δ-approximate α'-equilibrium, since otherwise it follows from Lemma C.2 and Definition 5.1 that the dynamics reached a $\beta(1+O(\delta))$-approximation of the optimal profit. We show that in each interval of T steps the potential increases by a factor of at least $\frac{\alpha\delta}{8\gamma n}$. Let $S^0, S^1, \ldots, S^T$ denote the joint actions of the players in times $0, 1, \ldots, T$ of this interval respectively. Since S^0 is not a δ-approximate α'-equilibrium, there exists a player with an improvement α'-move. Consider player j with the maximum absolute improvement α'-move and let S'_j be his best response. Recall that $\Delta_j(S^0) = p_j(S^0_{-j}, S'_j) - p_j(S^0)$. Let $\Delta'_j = \Delta_j(S^0)$ and let t' be the first time in this interval that player j is allowed to move. We denote by U the set of times before time t', where players made a best-response α-moves and we denote by $w(t)$ the player that moved at time t for each $t \in U$. Let $A = \sum_{t\in U} p_{w(t)}(S^{t+1})$ be the sum of the profits of the moving players after they make their moves. Now, we consider two cases:
Case 1: $A \leq \frac{\Delta'_j}{4\gamma}$. By the first property of the bounded jump assumption we have

$$p_j(S^{t'}) \leq p_j(S^0) + A \leq p_j(S^0) + \frac{\Delta'_j}{4\gamma} \leq p_j(S^0) + \frac{\Delta'_j}{4}. \tag{6}$$

By the second property of the bounded jump assumption we have

$$p_j(S^{t'}_{-j}, S'_j) \geq p_j(S^0_{-j}, S'_j) - \gamma \cdot A \geq p_j(S^0_{-j}, S'_j) - \gamma\frac{\Delta'_j}{4\gamma} = p_j(S^0_{-j}, S'_j) - \frac{\Delta'_j}{4}.$$

Hence,

$$p_j(S^{t'}_{-j}, S'_j) - p_j(S^{t'}) \geq p_j(S^0_{-j}, S'_j) - p_j(S^0) - \frac{\Delta'_j}{2} \geq \Delta'_j - \frac{\Delta'_j}{2} = \frac{\Delta'_j}{2}. \tag{7}$$

By equation (6) we get

$$p_j(S^{t'}) \leq p_j(S^0) + \frac{\Delta'_j}{4} < \frac{\Delta'_j}{\alpha'} + \frac{\Delta'_j}{4} < 2\frac{\Delta'_j}{4\alpha} = \frac{\Delta'_j}{2\alpha}.$$

Where the second inequality follows from the fact that Δ'_j is the improvement of player j when making his best response, which is an α'-move in step 0. Thus,

$$\alpha \cdot p_j(S^{t'}) < \frac{\Delta'_j}{2}. \tag{8}$$

By inequalities (7) and (8) we have

$$\alpha \cdot p_j(S^{t'}) < p_j(S^{t'}_{-j}, S'_j) - p_j(S^{t'})$$

Therefore, player j can make a best response α-move at time t' and increase the potential ϕ by at least $\frac{\Delta'_j}{2} \geq \frac{\delta}{2n}\phi(S^0)$.

Case 2: $A > \frac{\Delta'_j}{4\gamma}$. Since A is the sum of the profits of players making a best-response α-move after making the move, these players increase the potential ϕ by at least $\frac{\alpha}{1+\alpha}A > \frac{\alpha}{1+\alpha} \cdot \frac{\Delta'_j}{4\gamma} \geq \frac{\alpha\Delta'_j}{8\gamma} \geq \frac{\alpha\delta}{8\gamma n}\phi(S^0)$.

Let $\phi(i)$ denote the potential in round i. Then, in both cases $\phi(i) \geq \phi(S_{init})(1 + \frac{\alpha\delta}{8\gamma n})^i$. Since $\phi(i) \leq \phi^*$, the upper bound on the number of steps follows. ■

Remark C.5 The above theorem shows that we reach a state with profit at least $\frac{1}{\beta(1+O(\delta))}$ of the optimum after a polynomial number of best-response α-moves. Eventhough after this state, the profit of solutions can decrease, it follows from the proof of the theorem that the number of states in which the profit of the solution is less than $\beta(1 + O(\delta))$-approximation is at most $O(\frac{\gamma n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*}) \cdot T)$. In addition, since the potential function is always increasing after any α-move, the profit can decrease by a factor of at most $\frac{profit(S)}{\phi(S)}$. It is not hard to show that the ratio $\frac{profit(S)}{\phi(S)}$ for any strategy profile in cut games and market sharing games are exactly 1 and at most $\log(n)$ respectively. As a result, for cut games, the profit of any state after a polynomial number of steps reach a $2+O(\delta)$-approximate solution and remains within this factor afterwards. Also for market sharing games, the profit of any state after a polynomial number of steps reach a $2+O(\delta)$-approximate solution and remains within a factor of $O(\log n)$ of the optimal solution.

D Cut Games

In this section, we study convergence in cut games (also called the party affiliation games). We show that these games are *nice* games that satisfy the bounded jump condition. First, we show that cut games are 2-*nice* according to definition 5.1.

Lemma D.1 *Cut games are β-*nice *potential games with $\beta = 2$.*

Proof: We need to show that for any strategy profile S, $2(profit(S) + \Delta(S)) \geq OPT$. To do so, we show that $2(profit(S) + \Delta(S)) \geq \sum_{v \in V(G)} w_v$. Given any strategy profile S, for any player v, either $p_v(S) > \frac{w_v}{2}$, or if $p_v(S) < \frac{w_v}{2}$, then $\Delta_v(S) \geq w_v - p_v(S) - p_v(S)$, thus $2(p_v(S)+\Delta(S)) \geq 2(w_v - p_v(S)) \geq 2(w_v - \frac{w_v}{2}) = w_v$. Therefore, the cut game is a 2-nice game. ■

Lemma D.2 *Cut games satisfy the* 1*-bounded jump property.*

Proof: For two players u and v, if player u changes his strategy and goes to the same side as v, then payoff of v does not increase at all, thus $p_v(S_{-u}, S'_u) \leq p_v(S) + p_u(S_{-u}, S'_u)$. Otherwise, if player u changes his strategy to the other side of player v, the increase in the payoff of player v is at most $w_{u,v}$. Thus, $p_v(S_{-u}, S'_u) \leq p_v(S) + w_{u,v} \leq p_v(S) + p_u(S_{-u}, S'_u)$. This implies the first condition of the bounded jump property.

Now, consider a strategy profile S and two players u and v with two new strategies S'_u and S'_v. When player u changes his strategy to S'_u, if he decreases the payoff of strategy S'_v for player v, then it decreases this payoff by at most $w_{u,v}$. In this case, the payoff of u from switching to his strategy is at least $w_{u,v}$, therefore, $p_v(S_{-\{u,v\}}, S'_u, S'_v) \geq p_v(S_{-v}, S'_v) - p_u(S_{-u}, S'_u)$ which is the second condition of the bounded jump property.

■

Corollary D.3 *Let $\frac{1}{8} > \delta \geq 4\alpha$. Consider a cut game Λ with and any initial state S_{init}. The unrestricted α-Nash best-response dynamics with a liveness property generates a profile S with profit at least $\frac{1}{(2+O(\delta))} OPT(\Lambda)$ in at most$O\left(\frac{n}{\alpha\delta}\log(\frac{\phi(S_{init})}{\phi^*}) \cdot T\right)$ steps.*

E Market Sharing Games

In this section we consider market sharing games. We show that these games are 2-*nice* games that satisfy the 1-bounded jump condition. First, we show that congestion games with linear latency functions are 2-*nice* according to definition 5.1.

Lemma E.1 *Market sharing games are β-*nice *potential games with $\beta = 2$.*

Proof: We need to show that for any strategy profile S, $2(profit(S) + \Delta) \geq OPT$. To do so, we can show that $profit(S) + \sum_{i \in N} p_i(S_{-i}, S'_i) \geq OPT$ where S'_i is the best response of player i in strategy profile S. Let S^* be the strategy profile of the optimal solution. Then $p_i(S_{-i}, S'_i) \geq p_i(S_{-i}, S^*_i)$. Let T be the set of markets that are satified in the optimal solution , i.e., $OPT = \sum_{j \in T} v_j$. Let R be the set of markets in T that are satisfied in S and L be the rest of markets in T. All of markets in R are satisfied in S, thus the sum of profits of markets in R is less than $profit(S)$. Moreover, for any market j in L, if $j \in S^*_i$, then the profit $p_i(S_{-i}, S^*_i)$ contains the whole value v_j of market j, since no other player plays this market. Therefore, $\sum_{j \in L} v_j \leq \sum_{i \in N} p_i(S_{-i}, S^*_i) \leq \sum_{i \in N} p_i(S_{-i}, S'_i)$. The above inequalities imply the 2-nice property as follows:

$$OPT = \sum_{j \in T} v_j = \sum_{j \in R} v_j + \sum_{j \in L} v_j \leq profit(S) + \sum_{i \in N} p_i(S_{-i}, S'_i) \leq 2(profit(S) + \Delta).$$

■

Lemma E.2 *Market sharing games satisfy the* 1*-bounded jump property.*

Proof: Consider two players i and i' in strategy profile S. If player i' changes his best response strategy to $S'_{i'}$, the congestion of each market j changes from vector n_j to n'_j where $n_j - 1 \leq n'_j \leq n_j + 1$. Then the increase in the payoff of player i is at most $\sum_{j \in S_i \cap (S_{i'} \backslash S'_{i'})} (\frac{v_j}{n_j - 1} - \frac{v_j}{n_j})$. The payoff of player i' after changing his strategy from S_i to $S'_{i'}$ is at least $\sum_{j \in S_{i'}} \frac{v_j}{n_j}$. For a market $j \in S_i \cap (S_{i'} \backslash S'_{i'})$, at least two players i and i' are playing market j in S, thus $n_j \geq 2$, thus $(\frac{v_j}{n_j - 1} - \frac{v_j}{n_j}) \leq \frac{v_j}{n_j}$. Therefore,

$$\sum_{j \in S_i \cap (S_{i'} \backslash S'_{i'})} (\frac{v_j}{n_j - 1} - \frac{v_j}{n_j}) \leq \sum_{j \in S_i \cap (S_{i'} \backslash S'_{i'})} \frac{v_j}{n_j} \leq \sum_{j \in S_{i'}} \frac{v_j}{n_j} = p_{i'}(S).$$

This implies the first condition of the bounded jump property, i.e, the increase in the payoff of player i is at most the payoff i'.

Consider a strategy profile S and two players i and i' with two best response strategies S'_i and $S'_{i'}$. When player i' changes his strategy to $S'_{i'}$, if he decreases the payoff of strategy S'_i for player i, then it decreases this payoff by at most $\sum_{j \in S'_i \cap (S'_{i'} \backslash S_{i'})} (\frac{v_j}{n_j} - \frac{v_j}{n_j + 1})$. In this case, the payoff of i' from switching to his strategy is at least $\sum_{j \in S'_{i'}} \frac{v_j}{n_j + 1}$. Since for any market $j \in S'_i \cap (S'_{i'} \backslash S_{i'})$, we have $n_j \geq 1$, thus, $\frac{v_j}{n_j + 1} \geq \frac{v_j}{n_j} - \frac{v_j}{n_j + 1}$. These inequalities imply the second condition of the 1-bounded jump property as follows:

$$p_i(S_{-i}, S'_i) - p_i(S_{-\{i,i'\}}, S'_i, S'_{i'}) \leq \sum_{j \in S'_i \cap (S'_{i'} \backslash S_{i'})} (\frac{v_j}{n_j} - \frac{v_j}{n_j + 1}) \leq \sum_{j \in S'_{i'}} \frac{v_j}{n_j + 1} \leq p_{i'}(S_{-i'}, S'_{i'}).$$

■

Corollary E.3 *Let* $\frac{1}{8} > \delta \geq 4\alpha$*. Consider a market sharing game* Λ *with and any initial state* S_{init}*. The unrestricted* α*–Nash best-response dynamics with liveness property generates a profile* S *with profit* $\frac{1}{(2+O(\delta))} OPT(\Lambda)$ *in at most* $O\left(\frac{n}{\alpha\delta} \log(\frac{\phi^*}{\phi(S_{init})}) \cdot T\right)$ *steps.*