

Atomic congestion games: fast, myopic *and* concurrent*

Dimitris Fotakis³, Alexis C. Kaporis^{1,2}, and Paul G. Spirakis^{1,2}

¹ Dept. of Computer Eng. and Informatics, Univ of Patras, 26500 Patras, Greece

² Research Academic Comp. Tech. Inst., N. Kazantzaki Str, 26500 Patras, Greece

³ Dept. of Information & Communication Systems Eng., Univ. of the Aegean, Samos, Greece
Email: fotakis@aegean.gr, kaporis@ceid.upatras.gr, spirakis@cti.gr

Abstract. We study here the effect of concurrent greedy moves of players in atomic congestion games where n selfish agents (players) wish to select a resource each (out of m resources) so that her selfish delay there is not much. The problem of “maintaining” global progress while allowing concurrent play is exactly what is examined and answered here. We examine two orthogonal settings : (i) A game where the players decide their moves without global information, each acting “freely” by sampling resources randomly and locally deciding to migrate (if the new resource is better) via a random experiment. Here, the resources can have quite arbitrary latency that is load dependent. (ii) An “organised” setting where the players are pre-partitioned into selfish groups (coalitions) and where each coalition does an improving coalitional move. Our work considers concurrent selfish play for arbitrary latencies for the first time. Also, this is the first time where fast coalitional convergence to an approximate equilibrium is shown.

1 Introduction

Congestion games (CG) provide a natural model for non-cooperative resource allocation and have been the subject of intensive research in algorithmic game theory. A *congestion game* is a non-cooperative game where selfish players compete over a set of resources. The players’ strategies are subsets of resources. The cost of each player from selecting a particular resource is given by a non-negative and non-decreasing latency function of the load (or congestion) of the resource. The individual cost of a player is equal to the total cost for the resources in her strategy. A natural solution concept is that of a pure Nash equilibrium (NE), a state where no player can decrease his individual cost by unilaterally changing his strategy. In a classical paper, Rosenthal [27] showed that pure Nash equilibria on atomic congestion games correspond to local minima of a natural potential function. Twenty years later, Monderer and Shapley [24] proved that congestion games are equivalent to potential games. Many recent contributions have provided considerable insight into the structure and efficiency (e.g. [14, 3, 7, 17]) and tractability [12, 1] of NE in congestion games. Given the non-cooperative

* The 2nd and 3rd author were partially supported by the IST Program of the European Union under contract number IST-015964 (AEOLUS). This work was partially supported by the Future and Emerging Technologies Unit of EC (IST priority – 6th FP), under contract no. FP6-021235-2 (project ARRIVAL).

nature of congestion games, a natural question is whether the players trying to improve their cost converge to a pure NE in a reasonable number of steps. The potential function of Rosenthal [27] decreases every time a *single* player changes her strategy and improves her individual cost, while this is not true when concurrent selfish moves are performed. Hence every sequence of improving moves will eventually converge to a pure Nash equilibrium. However, this may require an exponential number of steps, since the problem is *PLS-complete* [12]. A pure Nash equilibrium of a *symmetric network* atomic congestion game can be found by a min-cost flow computation [12]. Even better, for *singleton* CG (aka CG on parallel links), for CG with *independent resources*, and for *matroid* CG, every sequence of improving moves reaches a pure Nash equilibrium in a polynomial number of steps [21, 1]. An alternative approach to circumvent the PLS-completeness of computing a pure Nash equilibrium is to seek an *approximate* NE. [6] considers *symmetric* congestion games with a weak restriction on latency functions and proves that several natural families of ε -moves converge to an ε -NE in time polynomial in n and ε^{-1} . However, sequential moves take $\Omega(n)$ steps in the worst case to reach an (approximate) NE and requires central coordination. A natural question is whether concurrent and autonomous play can convergence to an approximate pure Nash equilibrium. In this work, we investigate the effect of concurrent moves on the rate of convergence to approximate pure Nash equilibria.

1.1 Singleton Games with Myopic Players.

Related Work and Motivation. The Elementary Step System hypothesis, under which at most one user performs an improving move in each round, greatly facilitates the analysis of [8, 11, 17, 18, 22, 23, 26]. This is not an appealing scenario to modern networking, where simple decentralized distributed protocols can reflect better the essence of net’s liberal nature on decision making. All the above manifest the importance of distributed protocols that allow an arbitrary number of users to reroute per round, on the basis of selfish migration criteria. This is an Evolutionary Game Theory [31] perspective, see also [28] with a current treatment of both nonatomic games and of evolutionary dynamics. In this setting, the main concern is on studying the *replicator-dynamics*, that is to model the way that users revise their strategies throughout the process.

Discrete setting. The work in [10] considers n players concurrently sample for a better link amongst m parallel links per round (singleton CG). Link j has linear latency $s_j x_j$, where x_j is the number of players and s_j is the constant speed of the link j . This migration protocol uses global info: only users with latency exceeding the overall average link latency \bar{L}_t at round t are allowed with an appropriate probability to sample for a new link j . Also global info is used to amplify favorable links: link j is sampled *proportionally* to $d_t(j) = n_t(j) - s_j \bar{L}_t$, where $n_t(j)$ is the number of users on link j , and reaches in expectedly $O(\log \log n + \log m)$ rounds a NE. In [4] it was given the analysis of a concurrent protocol on identical links and players. On parallel during round t , each user b on resource i_b with load $X_{i_b}(t)$ selects a random resource j_b and if $X_{i_b}(t) > X_{j_b}(t)$ then b migrates to j_b with probability $1 - X_{j_b}(t)/X_{i_b}(t)$. It reaches an ε -NE in $O(\log \log n)$, or an exact NE in $O(\log \log n + m^4)$ rounds, in expectation.

Continuous setting. The work in [29] gives a general definition of nonatomic potential games, and shows convergence to Nash equilibrium in these games, under a very broad

class of evolutionary dynamics. A series of papers [5, 13] on the *Wardrop* model give strong intuition on this subject. In [13] the significance of the *relative slope* parameter d is shown. A latency function ℓ has relative slope d if $x\ell'(x) \leq d\ell(x)$. Each user on path P in commodity i , either with probability β selects a uniformly random path Q in i , or with probability $1 - \beta$ selects a path Q with probability proportional to its flow f_Q . If $\ell_Q < \ell_P$ user migrates to sampled Q with probability $\frac{\ell_P - \ell_Q}{d(\ell_P + \alpha)}$, where parameter α is arbitrary. In [5] it was shown that as long as all players concurrently employ arbitrary *no-regret* policies, they will eventually achieve convergence.

Contribution. We study a simple distributed protocol for congestion games on parallel links under very general assumptions on the latency functions. In parallel each player selects a link uniformly at random in each round and checks whether she can significantly decrease her latency by moving to the chosen link. If this is the case, the player becomes a potential migrant. The protocol selects at most one potential migrant to defect from each link. This is a local decision amongst users on the same link, allowing a realistic amount of parallelism amongst entities on different resources. Details on this, falling in the context of *dimension-exchange* protocols on load balancing, can be found in [2, 9, 16, 20]. We prove that if the number of players is $\Theta(m)$, the protocol reaches an almost-NE in $O(\log(\Phi_0/\Phi^*))$ time, where Φ_0 is Rosenthal's potential value as the game starts and Φ^* is the corresponding value at a NE. The proof of convergence is technically involved and interesting and comprises the main technical merit of this work. Our notion of approximate pure Nash equilibrium, see Definition 2, is a bit different from similar approximate notions considered in previous work [6, 10] in an atomic setting, while it is close in nature to the stable state defined in [13, Def. 4] for the Wardrop model. An *almost-Nash equilibrium* is a state where at most $o(m)$ links have latency either considerably larger or considerably smaller than the current average latency. This definition relaxes the notion of exact pure NE and introduces a meaningful notion of approximate (bicriteria) NE for our fully myopic model of migration described above. In particular, an almost-NE guarantees that unless a player uses an overloaded link (i.e. a link with latency considerably larger than the average latency), the probability that she finds (by uniform sampling) a link to migrate and significantly improve her latency is at most $o(1)$. Furthermore, it is unlikely that the almost-NE reached by our protocol assigns any number of players to overloaded an almost-NE). As it will become clear from the analysis, the reason that users do not accumulate on overloaded links, is that the number of players on such links is a strong super-martingale. In addition, by the fact that any bin initially has $O(\log n)$ load we get that in $O(\log n)$ rounds the overloaded bins will drain from users.

Our results extend the results in [4, 10] in the sense that (i) we consider arbitrary and unknown latency functions subject only to the α -bounded jump condition [6, Section 2], (ii) it requires no other global info. Also, the strategy space of player i may be extended to all subsets of resources of cardinality k_i such that $\sum_i k_i = O(m)$, see also independent resource CG [21].

1.2 Congestion Games with Coalitions

In many practical situations however, the competition for resources takes place among coalitions of players instead of individuals. For a typical example, one may consider

a telecommunication network where antagonistic service providers seek to minimize their operational costs while meeting their customers' demands. In this and many other natural examples, the number of coalitions (e.g. service providers) is rather small and essentially independent of the number of players (e.g. users). In addition, the coalitions can be regarded as having a quite accurate picture of the current state of the game and moving greedily and sequentially. In such settings, it is important to know how the competition among coalitions affects the rate of convergence to an (approximate) pure Nash equilibrium. Motivated by similar considerations, [19, 15] proposed *congestion games with coalitions* as a natural model for investigating the effects of non-cooperative resource allocation among static coalitions. In congestion games with coalitions, the coalitions are static and the selfish cost of each coalition is the total delay of its players. [19] mostly considers congestion games on parallel links with identical users and convex delays. For this class of games, [19] establishes the existence and tractability of pure NE, presents examples where coalition formation deteriorates the efficiency of NE, and bounds the efficiency loss due to coalition formation. [15] presents a potential function for linear congestion games with coalitions.

Contribution. In this setting, we present an upper bound on the rate of convergence to approximate pure Nash equilibria in single-commodity linear congestion games with static coalitions. The restriction to linear latencies is necessary because this is the only class of latency functions for which congestion games with static coalitions is known to admit a potential function and a pure Nash equilibrium. We consider ε -moves, i.e. deviations that improve the coalition's total delay by a factor more than ε . Combining the approach of [6] with the potential function of [15, Theorem 6], we show that if the coalition with the largest improvement moves in every round, an approximate NE is reached in a small number of steps. More precisely, we prove that for any initial configuration s_0 , every sequence of largest improvement ε -moves reaches an approximate NE in at most $\frac{kr(r+1)}{\varepsilon(1-\varepsilon)} \log \Phi(s_0)$ steps, where k is the number of coalitions, $r = \lceil \max_{j \in [k]} \{n_j\} / \min_{j \in [k]} \{n_j\} \rceil$ denotes the ratio between the size of the largest coalition and the size of the smallest coalition, and $\Phi(s_0)$ is the initial potential. This bound holds even for coalitions of different size, in which case the game is *not symmetric*. Since the recent results of [6] hold for symmetric games only, this is the first non-trivial upper bound on the convergence rate to approximate NE for a natural class of *asymmetric* congestion games. This bound implies that in *network* congestion games, where a coalition's best response can be computed in polynomial times by a min-cost flow computation [12, Theorem 2], an ε -Nash equilibrium can be computed in polynomial time. Moreover, in the special case that the number of coalitions is constant and the coalitions are almost equisized (i.e. $k = \Theta(1)$ and $r = \Theta(1)$), the number of ε -moves to reach an approximate NE is logarithmic in the initial potential.

2 Concurrent atomic congestion games

Model. There is a finite set of players $\{1, \dots, n\}$ and a set of edges (or resources) $E = \{e_1, \dots, e_m\}$. The strategy space S_i of player i is E . It is assumed that $n = O(m)$. The game consists of a sequence of rounds $t = 0, \dots, t^*$. It starts at round $t = 0$, where each player i selects myopically strategy $s_i(0) \in S_i$. In each subsequent round

$t = 1, \dots, t^*$, concurrently and independently, each player updates his current strategy $s_i(t)$ to $s_i(t+1)$ according to the simple, oblivious and distributed protocol `Greedy` presented in Section 2.1. That is, at round t the state $s(t) = \langle s_1(t), \dots, s_n(t) \rangle \in S_1 \times \dots \times S_n$ of the game is a combination of strategies over players. The number $f_e(t)$ of players on edge $e \in E$ is $f_e(t) = |\{j : e \in s_j(t)\}|$. Edge e has a latency $\ell_e(f_e(t))$ measuring the common delay of players on it at state $s(t)$. The cost $c_i(t)$ of player i equals the sum of latencies of all edges belonging in his current strategy $s_i(t)$, that is $c_i(t) = \sum_{e \in s_i(t)} \ell_e(f_e(t))$. Let the average delay of the resources be $\bar{\ell}(t) = \frac{1}{m} \sum_{e \in E} \ell_e(f_e(t))$. Consider the value of Rosenthal's potential $\Phi(t) = \sum_{e \in E} \sum_{x=1}^{f_e(t)} \ell_e(x)$. We assume no latency-info other than the α -bounded jump condition:

Definition 1. [6] Consider a set of m resources E each $e \in E$ incurring latency $\ell_e(x)$ when x players use it, $x \in \{0, \dots, n\}$. Let $\alpha = \min_a \{a \mid \forall x = 0, \dots, n, \forall e \in E \text{ it holds } \ell_e(x+1) \leq a \ell_e(x)\}$. Then each $e \in E$ satisfies the α -bounded jump condition.

This condition imposes a minor restriction on the increase-rate of the latency function $\ell_e(\cdot)$ of any resource $e \in E$. For example $\ell_e(x) = \alpha^x$ is α -bounded, which is also true for polynomials of degree $d \leq \alpha$. Our bicriterial equilibria (see [13, Def. 4]) follow.

Definition 2. An almost-NE is a state where $o(m)$ used edges have latency $> \alpha \bar{\ell}(t)$ and $\forall \epsilon > 0, \exists S \subseteq E : |S| \geq \epsilon m$ with used edges in S of latency $< \frac{1}{\alpha_S} \bar{\ell}(t)$, where α_S is the jump-parameter with respect to edges in S .

Target. We establish the following for protocol `Greedy` presented in Section 2.1.

Theorem 1. The expected number of rounds until `Greedy` reaches an almost-NE is at most $2 \lceil p^{-1} \ln(2\Phi_{\max}/\Phi_{\min}) \rceil$.

Constant $p = \Theta(1)$ is defined in Theorem 2, intuitively it provides a bound on the expected potential's drop caused by `Greedy` within all consecutive rounds which are not on an almost-NE. Theorem 1 follows easily (see the proof in the full version of the paper in [30]) from Theorem 2, see in turn its proof plan in Section 2.2. Here $\Phi_{\max}, (\Phi_{\min})$ denote the initial (final) value of the potential (value of the potential at an exact NE).

Taking into account the very limited info that our protocol extracts per round, our analysis suggests that an almost-NE of this kind is a meaningful notion of a stable state that can be reached quickly. In particular, the almost-NE reached by our protocol is a relaxation of an exact NE where the probability that a significant number of players can find (by uniform sampling) links to migrate and significantly improve their cost is small.

More precisely, in an exact NE, no used link has latency greater than $\alpha \bar{\ell}(t)$ and no link with positive load has latency less than $\bar{\ell}(t)/\alpha$, while the definition of an almost-NE imposes the same requirements on all but $o(m)$ links. Hence the notion of an almost-NE is a relaxation of the notion of an exact NE. In addition, a player not assigned to an overloaded link (i.e. a link with latency greater than $\alpha \bar{\ell}(t)$) can significantly decrease her cost (i.e. by a factor greater than α^2) only if she samples an underloaded link (i.e. a link with latency less than $\bar{\ell}(t)/\alpha$). Therefore, in an almost-NE, the probability that

a player not assigned to an overloaded link samples a link where she can migrate and significantly decrease her cost is $o(1)$. Furthermore, it is unlikely that the almost-NE reached by our protocol assigns a large number of players to overloaded links ⁴.

Theorem 2. *If round t is not an almost-NE then $\mathbb{E}[\Phi(t+1)] \leq (1-p)\mathbb{E}[\Phi(t)]$, with p bounded below by a positive constant.*

The proof plan of this theorem is presented in Section 2.2. Its proof will be given in Section 2.6 which combines results proved in Section 2.3, 2.4 and 2.5.

2.1 Concurrent protocol Greedy

Initialization: $\forall i \in \{1, \dots, n\}$ select a random $e \in \{1, \dots, m\}$.

During round t , do in parallel $\forall e \in E$:

1. Select 1 player i from e at random.
 2. Let player i sample for a destination edge e' u.a.r. over E .
 3. If $\ell_{e'}(f_{e'}(t))(\alpha + \delta_\vartheta) < \ell_e(f_e(t))$ then allow player i migrate to e' with probability $\vartheta = \Omega(1)$.
-

For $\vartheta, \delta_\vartheta$ see Section 2.3 Lemma 2, Corollary 1, and Section 2.5 Case 1 and 2.

2.2 Convergence of Greedy - Overview

The idea behind main Theorem 1 is to show that, starting from $\Phi(0) = \Phi_{\max}$, per round t of Greedy not in an almost-NE, the expected $\mathbb{E}[\Delta\Phi(t)]$ potential drop is a positive portion of the potential $\Phi(t)$ at hand. Since the minimum potential Φ_{\min} is a positive value, the total number of round is at most logarithmic in $\frac{\Phi_{\max}}{\Phi_{\min}}$. We present below how Sections 2.3, 2.4 and 2.5 will be combined together towards showing that Greedy gives a large “bite” to the potential $\mathbb{E}[\Phi(t)]$ at hand, per round not in an almost-NE, and prove key Theorem 2. Section 2.3 shows that $\mathbb{E}[\Delta\Phi(t)]$ is at most the total expected cost-drop $\sum_i \mathbb{E}[\Delta c_i(t)]$ of users allowed by Greedy to migrate and proves that $\sum_i \mathbb{E}[\Delta c_i(t)] < 0$, i.e. *super-martingale* [25, Def. 4.7]. Hence, showing large potential drop per round not in an almost-NE reduces to showing $\sum_i \mathbb{E}[\Delta c_i(t)]$ equals a positive number times $-\mathbb{E}[\Phi(t)]$. This is achieved in Sections 2.4 and 2.5 which show that $|\sum_i \mathbb{E}[\Delta c_i(t)]|$ and $\mathbb{E}[\Phi(t)]$ are both closely related to $\mathbb{E}[\bar{\ell}(t)] \times m$, i.e. both are a corresponding positive number times $\mathbb{E}[\bar{\ell}(t)] \times m$. First, Section 2.4 shows that $\mathbb{E}[\Phi(t)]$ is a portion of $\mathbb{E}[\bar{\ell}(t)] \times m$. Having this, fast convergence reduces to showing $\sum_i \mathbb{E}[\Delta c_i(t)]$ equals a positive number times $-\mathbb{E}[\bar{\ell}(t)] \times m$ which is left to Section 2.5 & 2.6. At the end, Section 2.6 puts together Sections 2.3, 2.4 and 2.5 and completes the proof of our key Theorem 2.

⁴ Due to the initial random allocation of the players to the links, the overloaded links (if any) receive $O(\log n)$ players with high probability. Lemma 3 and Corollary 2 show that the number of players on any overloaded link is a strong super-martingale during each round. Thus, such overloaded links will drain from users in expectedly $O(\log n)$ rounds.

2.3 Showing that $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)]$ upper bounds $\mathbb{E}[\Delta \Phi(t)]$

Let $\mathcal{A}(t)$ the migrants allowed in step (3) of Greedy in Section 2.1. Linearity of expectation by Lemma 1 yields $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \geq \mathbb{E}[\Delta \Phi(t)]$. $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] < 0$ follows by Lemma 2 and Corollary 1 below: user $i \in \mathcal{A}(t)$, by selfish criterion in step (3) of Greedy, decreases expectedly its cost if the latency on i 's departure link is $> (\alpha + \delta_\vartheta)$ times the latency on its destination. Here ϑ is the migration probability in step (3) of Greedy.

Lemma 1. $\sum_{i \in \mathcal{A}(t)} \Delta[c_i(t)] \geq \Delta[\Phi(t)]$. Equality holds if $\Delta[f_e(t)] \leq 1, \forall e \in E$.

Proof. See the proof in the full version of the paper in [30]. □

Lemma 2. For every positive constant δ , if migration probability ϑ of Greedy is at most $\min\{\frac{\delta}{\alpha(\alpha-1)}, 1\}$, the expected latency of a destination link e in the next round $t+1$ is:

$$\mathbb{E}[\ell_e(f_e(t+1))] \leq (1 + \delta/\alpha)\ell_e(f_e(t) + 1) \leq (\alpha + \delta)\ell_e(f_e(t))$$

Proof. See the proof in the full version of the paper in [30]. □

Corollary 1. $\mathbb{E}[\Delta c_i(t) | c_i(t)] \leq \ell_{e'}(f_{e'}(t))(\alpha + \delta_\vartheta) - c_i(t) < 0, \forall i \in \mathcal{A}(t)$ migrating $e \rightarrow e'$.

Proof. See the proof in the full version of the paper in [30]. □

2.4 Showing that $\mathbb{E}[\Phi(t)]$ is at most a portion of $\mathbb{E}[\bar{\ell}(t)] \times m$

By Greedy's initialization the load is Binomially distributed, thus at round $t = 0$ we easily get (see the full version in [30]):

$$\mathbb{E}[\bar{\ell}(0)] \leq e^{\alpha \frac{n}{m-1}} e^{-\frac{n}{m}} = O(1), \text{ and } \mathbb{E}[\Phi(0)] = O(\mathbb{E}[\bar{\ell}(0)] \times m), \quad (1)$$

However, Greedy may affect badly the initial distribution of bins, thus making (1) invalid for each $t > 0$. We shall show that similar to round 0 strong tails will make (1) true for each round $t > 0$. To see this, consider the concurrent random process Blind (a simplification of Greedy in Section 2.1). At $t = 0$ throw randomly $n = O(m)$ balls to m bins (Blind's and Greedy's initializations are identical). Initially, the load distribution has Binomial tails from deviating from expectation $O(n/m) = O(1)$. During round $t > 0$, Blind draws exactly 1 random ball from each loaded bin (as Step 1 of Greedy). Let $n(t)$ the subset of drawn balls during round t . Round t ends by throwing at random these $|n(t)|$ drawn balls back into the m bins (then $|n(t)|$ allowed by Blind to migrate is at least the migrants allowed by Greedy, since no selfish criterion is required). Any bin is equally likely to receive any ball, thus, Blind preserves per round $t > 0$ strong Binomial tails from deviating from the constant expectation $O(n/m) = O(1)$ reminiscent to ones for $t = 0$. The above make true (1) for each round $t > 0$ of Blind.

Towards showing that Greedy also behaves, on a proper subset of bins, similarly to Blind it is useful the following lemma. Lemma 3 and Corollary 2 prove a supermartingale property on the load of bins with latency greater than a critical constant. This will help us to identify this subset of critical bins that will preserve similar bounds to (1) for each round $t > 0$ of Greedy.

Lemma 3. *Let ν be any integer no less than $\lceil 2n/m \rceil + 1$. For any round $t \geq 0$, every link e with $\ell_e(f_e(t)) \geq \alpha^\nu$ has $\mathbb{E}[f_e(t+1)] \leq f_e(t)$.*

Proof. See the proof in the full version of the paper in [30]. □

Corollary 2. *Consider the corresponding numbers ν 's defined in Lemma 3. We can find a constant $L^* : \forall t \geq 0$ on each edge with latency $\geq L^*$ the load is super-martingale.*

Let the constant L^* be as in Corollary 2 and define $\mathcal{A}_{L^*}(t) = \{e \in E : \ell_e(f_e(t)) < L^*\}$ and $\mathcal{B}_t^{L^*} = E \setminus \mathcal{A}_{L^*}(t)$. The target of Lemma 4 is to show that $\mathcal{B}_t^{L^*}$ is the subset of critical bins that will preserve similar bounds to (1) for each round $t > 0$ of Greedy.

Lemma 4.
$$\sum_{e \in \mathcal{B}_t^{L^*}} \frac{\mathbb{E}[\ell_e(f_e(t))]}{m} = O(1), \quad \sum_{e \in \mathcal{B}_t^{L^*}} \frac{\mathbb{E}[f_e(t)\ell_e(f_e(t))]}{m} = O(\mathbb{E}[\bar{\ell}(t)])$$

Proof. See the full version of the paper in [30]. □

Now, Fact 3 proves that $\mathbb{E}[\Phi(t)]$ is at most a portion of $\mathbb{E}[\bar{\ell}(t)] \times m$.

Fact 3 *If round t is not an almost-NE then $\mathbb{E}[\bar{\ell}(t)]m \geq \frac{\mathbb{E}[\Phi(t)]}{r(1+y_t)+1+x_t}$, $r = n/m$ and $r, y_t, x_t = \Theta(1)$.*

Proof. See the proof in the full version of the paper in [30]. □

2.5 Showing that $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)]$ is a portion of $-\bar{\ell}(t) \times m$

Sketch of Case 1 and 2 below. According to Definition 2, a round is not at an almost-NE if $\geq \varepsilon m$ links are either *overloaded* (of latency $\geq \alpha \times \bar{\ell}(t)$) or *underloaded* (of latency $\leq \frac{1}{\alpha} \times \bar{\ell}(t)$) ones. We study separately each of these options in Cases 1 and 2 below. In both cases we relate $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)]$ to $-\bar{\ell}(t) \times m$. The idea beyond both Case 1 and 2 is simple: each migrant from $\mathcal{O}(t)$ to $\mathcal{U}(t)$ will contribute to $\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)]$ her little portion of $-\bar{\ell}(t)$ at hand (by the martingale property on the expected gain per user $i \in \mathcal{A}(t)$ proved in Corollary 1 Section 2.3). It remains to show that such migrations have as high impact as to boost the tiny atomic gain of order $\bar{\ell}(t)$, when considered in the overall population of migrants $\mathcal{A}(t)$, up to a portion of $\bar{\ell}(t) \times m$. Towards this, Fact 4 and 5 below show that, as long as the state is not an almost-NE, it induces imbalance amongst link-costs, which in turn influences a sufficient amount of migrations as to get cost-drop of order $-\bar{\ell}(t) \times m$.

Case 1. Here we define *underloaded* links in round t be $\mathcal{U}(t) = \{e \in E : \ell_e(f_e(t)) < (1 - \delta)\bar{\ell}(t)\}$, while *overloaded* ones are $\mathcal{O}(t) = \{e \in E : \ell_e(f_e(t)) \geq \alpha\bar{\ell}(t)\}$. Let us assume that we are not at an almost-NE because $|\mathcal{O}(t)| \geq \varepsilon m$, with constant $\varepsilon \in (0, 1)$.

Fact 4 For every $\alpha > 1$ if $|\mathcal{O}(t)| \geq \varepsilon m$, then $|\mathcal{U}(t)| \geq \delta m$, with $\delta = \frac{\varepsilon}{2}(\alpha - 1)$.

Proof. See the proof in the full version of the paper in [30]. \square

Therefore, for every $e \in \mathcal{O}(t)$, a player migrates from e to a link in $\mathcal{U}(t)$ with probability at least $\vartheta\delta$ (see step (3) of Greedy, Section 2.1). Using Lemma 2 with $\vartheta = \varepsilon/4$, we obtain that the expected decrease in its cost is at least $\frac{\delta}{2}\alpha\bar{\ell}(t)$ (see the proof in the full version in [30]).

Given that k migrants switch from a link in $\mathcal{O}(t)$ to a link in $\mathcal{U}(t)$ we obtain that their expected cost-drop is at least $\frac{\delta}{2}\alpha\bar{\ell}(t)$ times their number k . Let $p_{\mathcal{O} \rightarrow \mathcal{U}}(k)$ the probability to have k such migrants. The expected number $\sum_k k p_{\mathcal{O} \rightarrow \mathcal{U}}(k)$ of such migrants is at least $\varepsilon\vartheta\delta m$, since for every $e \in \mathcal{O}(t)$ with $|\mathcal{O}(t)| \geq \varepsilon m$, exactly 1 player migrates from e to a link in $\mathcal{U}(t)$ with probability at least $\vartheta\delta$ (see Fact 4 and step (3) of Greedy, Section 2.1). Now, the unconditional on k expected cost-drop due to migrants switching from links in $\mathcal{O}(t)$ to links in $\mathcal{U}(t)$ is at least

$$\sum_k \left(\frac{\delta}{2}\alpha\bar{\ell}(t)k \times p_{\mathcal{O} \rightarrow \mathcal{U}}(k) \right) \geq \frac{\delta}{2}\alpha\bar{\ell}(t) \times \varepsilon\vartheta\delta m = \varepsilon\vartheta\frac{\delta^2}{2}\alpha m\bar{\ell}(t) \quad (2)$$

By (2) we finally prove (for Case 1) the result of this section:

$$\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq -\varepsilon\vartheta\frac{\delta^2}{2}\alpha \times \bar{\ell}(t)m \quad (3)$$

Case 2. Here we define as *underloaded* links in round t be $\mathcal{U}(t) = \{e \in E : \ell_e(f_e(t)) < \frac{1}{\alpha}\bar{\ell}(t)\}$ and *overloaded* ones in $\mathcal{O}(t) = \{e \in E : \ell_e(f_e(t)) \geq (1 + \delta)\bar{\ell}(t)\}$. Let us assume that we are not at an almost-NE because $|\mathcal{U}(t)| \geq \varepsilon m$.

Fact 5 If $|\mathcal{U}(t)| \geq \varepsilon m$, then $\sum_{e \in \mathcal{O}(t)} \ell_e(f_e(t)) > \delta\bar{\ell}(t)m$, with $\delta = \frac{\varepsilon(\alpha-1)}{2\alpha}$.

Proof. See the proof in the full version of the paper in [30]. \square

Since $|\mathcal{U}(t)| \geq \varepsilon m$, a player migrates from each $e \in \mathcal{O}(t)$ to a link in $\mathcal{U}(t)$ with probability at least $\vartheta\varepsilon$ (see step (3) of Greedy, Section 2.1). Using Lemma 2 with $\vartheta = \frac{\varepsilon}{4\alpha}$, we obtain that the expected decrease in the cost of such a player is at least $\frac{\delta}{2(1+\delta)}\ell_e(f_e(t)) \geq \frac{\delta}{4}\ell_e(f_e(t))$ (see the proof in the full version in [30]). Using Fact 5, we obtain that the expected cost-drop due to migrants leaving overloaded links $\mathcal{O}(t)$ and entering $\mathcal{U}(t)$ in round t is at least:

$$\vartheta\varepsilon \times \frac{\delta}{4} \sum_{e \in \mathcal{O}(t)} \ell_e(f_e(t)) > \vartheta\varepsilon \times \frac{\delta}{4} \times \delta\bar{\ell}(t)m > \frac{\vartheta\varepsilon\delta^2}{4}\bar{\ell}(t)m \quad (4)$$

By (4) we finally prove (for Case 2) the result of this section:

$$\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq -\frac{\vartheta\varepsilon\delta^2}{4} \times m\bar{\ell}(t) \quad (5)$$

2.6 Proof of key Theorem 2.

Here we combine the results in Section 2.3, 2.4 and 2.5 and prove Theorem 2. From Section 2.3 we get $\mathbb{E}[\Delta\Phi(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] < 0$. As long as Greedy does not reach an almost-NE because: (i) The *overloaded* links, with respect to the realization $\bar{\ell}(t)$, are $|\mathcal{O}(t)| \geq \varepsilon m$. Then, we get from Expression (3) in Section 2.5 that $\mathbb{E}[\Delta\Phi(t)|\bar{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\bar{\ell}(t)] < -\varepsilon\vartheta \frac{\vartheta^2}{2} \alpha \times \bar{\ell}(t)m$. (ii) The *underloaded* links, with respect to the realization $\bar{\ell}(t)$, are $|\mathcal{U}(t)| \geq \varepsilon m$. Then, we get from Expression (5) in Section 2.5 that $\mathbb{E}[\Delta\Phi(t)|\bar{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\bar{\ell}(t)] < -\frac{\vartheta\varepsilon\delta^2}{4} \times \bar{\ell}(t)m$. In either Case 1 or 2 such that an almost-NE is not reached by realization $\bar{\ell}(t)$, we conclude from the above:

$$\mathbb{E}[\Delta\Phi(t)|\bar{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\bar{\ell}(t)] < -\frac{\vartheta\varepsilon\delta^2}{4} \times \bar{\ell}(t)m \quad (6)$$

Consider the space of all realizations $\bar{\ell}(t)$ not in an almost-NE due to $\geq \varepsilon m$ overloaded or underloaded links in round t . Let $p_{\bar{\ell}(t)}$ the probability to obtain a realization $\bar{\ell}(t)$ in this space. Removing the conditional on $\bar{\ell}(t)$, Expression (6) becomes:

$$\begin{aligned} \mathbb{E}[\Delta\Phi(t)] &= \sum_{\bar{\ell}(t)} \mathbb{E}[\Delta\Phi(t)|\bar{\ell}(t)] p_{\bar{\ell}(t)} \leq \sum_{\bar{\ell}(t)} \left[\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\bar{\ell}(t)] \right] p_{\bar{\ell}(t)} \\ &\leq \sum_{\bar{\ell}(t)} \left[-\frac{\vartheta\varepsilon\delta^2}{4} \times \bar{\ell}(t)m \right] p_{\bar{\ell}(t)} = -\frac{\vartheta\varepsilon\delta^2}{4} \times \mathbb{E}[\bar{\ell}(t)]m \end{aligned}$$

From Fact 3 the above becomes: $\mathbb{E}[\Delta\Phi(t)] \leq -\frac{\vartheta\varepsilon\delta^2}{4} \times \frac{\mathbb{E}[\Phi(t)]}{r(1+y_t)+1+x_t}$, $r = n/m$ and $r, x_t, y_t = \Theta(1)$.

3 Approximate Equilibria in Congestion Games with Coalitions

3.1 Model and Preliminaries

A *congestion game with coalitions* consists of a set of identical players $N = [n]$ ($[n] \equiv \{1, \dots, n\}$) partitioned into k coalitions $\{C_1, \dots, C_k\}$, a set of resources $E = \{e_1, \dots, e_m\}$, a strategy space $\Sigma_i \subseteq 2^E$ for each player $i \in N$, and a non-negative and non-decreasing latency function $\ell_e : \mathbb{N} \mapsto \mathbb{N}$ associated with every resource e . In the following, we restrict our attention to games with linear latencies of the form $\ell_e(x) = a_e x + b_e$, $a_e, b_e \geq 0$, and symmetric strategies (or *single-commodity* congestion games), where all players share the same strategy space, denoted Σ . The congestion game is played among the coalitions instead of the individual players. We let n_j denote the number of players in coalition C_j . The strategy space of coalition C_j is Σ^{n_j} and the strategy space of the game is $\Sigma^{n_1} \times \dots \times \Sigma^{n_k}$. A pure strategy $s_j \in \Sigma^{n_j}$ determines a (pure) strategy $s_j^i \in \Sigma$ for every player $i \in C_j$. We should highlight that if the coalitions have different sizes, the game is *not symmetric*. We let $r \equiv \lceil \max_{j \in [k]} \{|C_j|\} / \min_{j \in [k]} \{|C_j|\} \rceil$ denote the ratio between the size of the largest

coalition to the size of the smallest coalition. Clearly, $1 \leq r < n$. For every resource $e \in E$, the load (or congestion) of e due to C_j in s_j is $f_e(s_j) = |\{i \in C_j : e \in s_j^i\}|$. A tuple $s = (s_1, \dots, s_k)$ consisting of a pure strategy $s_j \in \Sigma^{n_j}$ for every coalition C_j is a *state* of the game. For every resource $e \in E$, the load of e in s is $f_e(s) = \sum_{j=1}^k f_e(s_j)$. The delay of a strategy $\alpha \in \Sigma$ in state s is $\ell_\alpha(s) = \sum_{e \in \alpha} \ell_e(f_e(s))$. The selfish cost of each coalition C_j in state s is given by the *total delay* of its players, denoted $\tau_j(s)$. Formally, $\tau_j(s) \equiv \sum_{i \in C_j} \ell_{s_j^i}(s) = \sum_{e \in E} f_e(s_j) \ell_e(f_e(s))$. Computing a coalition's best response in a network congestion game can be performed by first applying a transformation similar to that in [12, Theorem 2] and then computing a min-cost flow. A state s is a *Nash equilibrium* if for every coalition C_j and every strategy $s_j' \in \Sigma^{n_j}$, $\tau_j(s) \leq \tau_j(s_{-j}, s_j')$, i.e. the total delay of coalition C_j cannot decrease by C_j 's unilaterally changing its strategy. For every $\varepsilon \in (0, 1)$, a state s is an ε -*Nash equilibrium* if for every coalition C_j and every strategy $s_j' \in \Sigma^{n_j}$, $(1 - \varepsilon)\tau_j(s) \leq \tau_j(s_{-j}, s_j')$. An ε -*move* of coalition C_j is a deviation from s_j to s_j' that decreases the total delay of C_j by more than $\varepsilon\tau_j(s)$. Clearly, a state s is an ε -Nash equilibrium iff no coalition has an ε -move available.

3.2 Convergence to Approximate Equilibria

To bound the convergence time to ε -Nash equilibria, we use the following potential function: $\Phi(s) = \frac{1}{2} \sum_{e \in E} [f_e(s)\ell_e(f_e(s)) + \sum_{j=1}^k f_e(s_j)\ell_e(f_e(s_j))]$, where [15, Theorem 6] proves that Φ is an exact potential function for (even multi-commodity) congestion games with static coalitions and *linear* latencies. We prove that for single-commodity linear congestion games with coalitions, the *largest improvement ε -Nash dynamics* converges to an ε -Nash equilibrium in a polynomial number of steps. Hence in network congestion games, where a coalition's best response can be computed in polynomial times by a min-cost flow computation, an ε -Nash equilibrium can be computed in polynomial time. If the current strategies profile is not an ε -Nash equilibrium, there may be many coalitions with ε -moves available. In the largest improvement ε -Nash dynamics, the coalition that moves is the one whose best response is an ε -move and results in the largest improvement in its total delay (and consequently in the potential). In the full version of the paper [30], the following theorem is proven.

Theorem 6. *In a single-commodity linear congestion game with n players divided into k coalitions, the largest improvement ε -Nash dynamics starting from an initial state s_0 reaches an ε -Nash equilibrium in at most $\frac{kr(r+1)}{\varepsilon(1-\varepsilon)} \log \Phi(s_0)$ steps, where $r = \lceil \max_{j \in [k]} \{n_j\} / \min_{j \in [k]} \{n_j\} \rceil$ denotes the ratio between the size of the largest coalition and the size of the smallest coalition.*

References

1. H. Ackermann, H. Roeglin, and B. Voeking. On the impact of combinatorial structure on congestion games. In *FOCS*, 2006.
2. H. Arndt. Load balancing: dimension exchange on product graphs. In *18th International Symposium on Parallel and Distributed Processing*, 2004.

3. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *STOC*, 2005.
4. P. Berenbrink, T. Friedetzky, L. A. Goldberg, P. Goldberg, Z. Hu, and R. Martin. Distributed selfish load balancing. In *SODA*, 2006.
5. A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: on convergence to nash equilibria of regret-minimizing algorithms in routing games. In *PODC*, 2006.
6. S. Chien and A. Sinclair. Convergence to Approximate Nash Equilibria in Congestion Games. In *SODA*, 2007.
7. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *STOC*, 2005.
8. G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *STACS*, 2006.
9. G. Cybenko. Load balancing for distributed memory multiprocessors. *J. Parallel Distrib. Comput.*, 7:279–301, 1989.
10. E. Even-Dar and Y. Mansour. Fast convergence of selfish rerouting. In *SODA*, 2005.
11. Eyal Even-Dar, Alex Kesselman, and Yishay Mansour. Convergence Time to Nash Equilibria. In *ICALP*, 2003.
12. A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *STOC*, 2004.
13. S. Fischer, H. Räcke, and B. Vöcking. Fast convergence to wardrop equilibria by adaptive sampling methods. In *STOC*, 2006.
14. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *TCS*, 348:226–239, 2005.
15. D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic Congestion Games among Coalitions. In *ICALP*, 2006.
16. B. Ghosh and S. Muthukrishnan. Dynamic load balancing in parallel and distributed networks by random matchings. In *SPAA*, 1994.
17. M. X. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS'05*.
18. P. W. Goldberg. Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In *PODC*, 2004.
19. A. Hayrapetyan, É. Tardos, and T. Wexler. The Effect of Collusion in Congestion Games. In *STOC*, 2006.
20. S. H. Hosseini, B. Litow, M. I. Malkawi, J. McPherson, and K. Vairavan. Analysis of a graph coloring based distributed load balancing algorithm. *J. Par. Distr. Comp.*, 10:160–166, 1990.
21. S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *AAAI*, 2005.
22. Lavy Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.
23. V. Mirrokni and A. Vetta. Convergence issues in competitive games. In *APPROX'04*.
24. D. Monderer and L. Shapley. Potential Games. *Games & Econ. Behavior*, 14:124–143, 1996.
25. R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
26. A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. on Net.*, 1(5):510–521, 1993.
27. R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
28. W. H. Sandholm. *Population Games and Evolutionary Dynamics*. To be published by MIT Press. <http://www.ssc.wisc.edu/whs/book/index.html>.
29. W. H. Sandholm. Potential Games with Continuous Player Sets. *Journal of Economic Theory*, 97:81–108, 2001.
30. Extended version. <http://students.ceid.upatras.gr/kaporis/papers/fks-sagt-08.pdf>.
31. J. W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.