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

Dimitris Fotakis[3], Alexis C. Kaporis[1], and Paul G. Spirakis[1,2]

[1] Dept of Computer Eng and Informatics, University of Patras, 26500 Patras, Greece
[2] Research Academic Computer Technology Institute, N. Kazantzaki Str, Patras University Campus, 26500 Patras, Greece
[3] Department of Information and Communication Systems Engineering, University of the Aegean
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. Such games usually admit a global potential that decreases by sequential and selfishly improving moves. However, concurrent moves may not always lead to global convergence. On the other hand, concurrent play is desirable because it might essentially improve the system convergence time to some balanced state. 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. Here the concurrency is among the members of the coalition. In this second setting, the resources have latency functions that are only linearly dependent on the load, since this is the only case so far where a global potential exists.
In both cases (i), (ii) we show that the system converges to an "approximate" equilibrium very fast (in logarithmic rounds where the logarithm is taken on the maximum value of the global potential). This is interesting, since two quite orthogonal settings lead to the same result. 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. All our results refer to atomic games (ie players are finite and distinct).

Keywords: Algorithms, concurrent atomic congestion games, Nash equilibria

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 [32] showed that pure Nash equilibria on atomic congestion games correspond to local minima of a natural potential function. Twenty years later, Monderer and Shapley [29] proved that congestion games are equivalent to potential games. Many recent contributions have provided considerable insight into the structure and efficiency (e.g. [16, 2, 8, 18]) and tractability [13, 1] of NE in congestion games.

Given the non-cooperative 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 [32] decreases every time a single player changes her strategy and improves her individual cost. Hence every sequence of improving moves will eventually converge to a pure Nash

⋆ The 2nd and 3rd author were partially supported by the IST Program of the European Union under contract number IST-015964 (AEOLUS).

equilibrium. However, this may require an exponential number of steps, since computing a pure Nash equilibrium of a congestion game is *PLS-complete* [13].

Nevertheless, there are many interesting classes of congestion games for which a pure Nash equilibrium can be computed in polynomial time. For example, a pure Nash equilibrium of a *symmetric network* atomic congestion game can be found by a min-cost flow computation [13]. 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 equilbrium in a polynomial number of steps [22, 1]. An alternative approach to circumvent the PLS-completeness of computing a pure Nash equilibrium is to seek an *approximate* NE (formally, an ε-NE), where no player can improve her cost by a factor more than ε by unilaterally changing her strategy. [7] 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, every family of sequential moves takes $\Omega(n)$ steps in the worst case to reach an (approximate) NE and its implementation requires central coordination among the players. In the view of the facts that the number of players is usually quite large and that central coordination between them is difficult to achieve, a natural question is whether concurrent play can accelerate the 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. Our main results concern two natural (and essentially orthogonal) settings where the rate of convergence is quite fast and mostly determined by the logarithm of the initial potential value.

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 [9, 11, 18, 19, 26, 27, 30]. However, a significant drawback of playing sequentially is that it requires $\Omega(n)$ rounds in the worst-case until n users reach a NE, not to mention the negative result [13] that holds on an atomic setting. Also, central control is imposed on moves. This is not an appealing scenario to modern networking, where simple decentralized distributed protocols can reflect better the essence of net's liberal nature. Furthermore, classical proofs of sequential convergence are based on assumptions of unbounded rationality and global knowledge. In real-world networks it is unrealistic to assume any player capable of monitoring the entire network per round. But even if a user can grasp the whole picture, it is computationally demanding to decide her best move.

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. It is important that migration rules are simple and myopic, while strong enough for the players to quickly reach (learn) a stable state. Here, terms "simple" and "myopic" mean that any selfish decision is taken by easy computations based on local info only, that is, the decision does not rely on global or expensive information about the overall current state of resources.

This is an Evolutionary Game Theory [33] perspective, which studies conditions under which a population of agents may (or may not) reach stable states. 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. Each user may revise her strategy by performing action *sampling* (a new resource is drawn) or *migration* (a move to a new resource). Sampling is further categorized as *uniform* (all resources are equally likely) or *proportional* (each resource is selected with probability proportional to a parameter related to it, usually, but not restricted to, the number of users on it). Uniform sampling is the cheapest way of searching the available resources. However, it typically results to slow convergence time, since it does not amplify highly appealing resources. On the contrary, proportional sampling, highly boosts the speed of the process, since it injects vast amounts of users into most appealing resources at hand. A word of caution, however, is that if the only sampling available is proportional to the number of users per resource, then the process becomes trapped only to the loaded resources up to now. A way out is to shift at an appropriate rate to uniform sampling, capable of exploring even currently empty resources.

At this point we should stress that, unlike sequential moves, the lack of global info and the fact that costs over resources may increase unboundedly on demand, it is possible that concurrent migrations oscillate the game eternally away from NE. Intuitively, while user i finds appealing a given resource e, simultaneously many other users may opt for e, increasing e's latency in a cost that ruins the profit of user i. This is a major difficulty on proving concurrent convergence. Such bad oscillation effects are known to Network and Telecommunications Community [23, 25, 31].

Let us first focus to the discrete concurrent setting. The work in [12] 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 is the KP model [24]. This migration protocol, although concurrent, is not completely decentralized, since it uses global info in order to allow only proper subsets of users to migrate. More precisely, on parallel, only users with latency exceeding the overall average link latency $\overline{L}_t$ at round t are allowed with an appropriate probability to sample for a new link j. We stress here that, for the case of multiple different links, this sampling for a link j is *proportional* to $d_t(j) = n_t(j) - s_j\overline{L}_t$, where $n_t(j)$ is the number of users on link j. Once more, this type of proportional sampling exposes global info to amplify favorable links, in contrast to the myopic scenario of sampling a random user, which in turn amplifies links proportionally to their load. All in all, these criteria highly boost the convergence time, requiring expectedly $O(\log\log n + \log m)$ rounds. On an experimental view, the work in [20] was prior to [12], were a series of similar concurrent protocols were validated.

In [4] it was given the analysis of a concurrent protocol on identical links and players. Notice here that the parallel links are identical, while the ones in [12] were related, but the important aspect of the analysis in [4] is that no global information was given to the migrants. 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 $1 - X_{j_b}(t)/X_{i_b}(t)$. Despite that users perform only uniform sampling, this protocol quickly reaches an ε-NE in $O(\log\log n)$, or an exact NE in $O(\log\log n + m^4)$ rounds, in expectation.

The reason that proportional sampling turns out to be not so crucial here, is the fact that all links are identical, so there is no need to inject many users to any particular speedy link. Thus, an important question is to what extent such myopic distributed protocols can cope with links that have large discrepancies amongst their latency functions.

Finally, we focus on the continuous concurrent setting. Powerful concurrent protocols have been analyzed in a continuous setting with respect to the Wardrop model (nonatomic flows) on general k commodities nets. The fact that each agent controls an infinitesimal amount of flow facilitates the analysis, since any concurrent migration of a lower order population of players causes almost no oscillation effect. However, a great difficulty that in turn arises here, is when a significant order of the population concurrently migrates. A series of important papers [5, 14] provide strong intuition on this subject. More precisely [14] shows the significance of the *relative slope* parameter d for the replicator dynamics to eventually converge to a stable state. Intuitively, a latency function ℓ has relative slope d if $x\ell'(x) \leq d\ell(x)$. Thus, parameter d is a peak-measure of ℓ and convergence can not occur if a link latency grows arbitrarily large with respect to flow fluctuations on it. The replication dynamics studied in [14] employs both uniform and proportional sampling. On parallel 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. Then, 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. However, probability β is rather cumbersome to tune, since it uses extensive information that concerns all latency functions and their corresponding first derivatives: $\beta \leq \frac{\min_{P\in\mathcal{P}} \ell_P(0)+\alpha}{L \max_{e\in E} \max_{x\in[0,\beta]} \ell'_e(x)}$.

While the work in [14] studies specific replication policies designated to yield fast convergence, in [5] it was shown a more general result. It stated that as along as all players concurrently employ arbitrary *no-regret* policies, they will eventually achieve convergence. Quoting from [5]: “any no-regret algorithm have the property that in any online, repeated game setting, their average loss per time step approaches that of the best fixed strategy in hindsight (or better) over time".

The work in [28, 15] remove the assumption of perfect information. In the sense that decisions are taken on the basis of a bulletin board which does not depict the most "fresh" state. If the info depicted on this board is too old and not regularly updated then oscillations occur. The analysis tunes the rate of updating the bulletin for the system eventually to convergence, see also [6, 10, 3]. More precisely, in [15] an arbitrary k commodity network is given with edge latency functions. Each user, independently and according to a Poisson distribution decides to revise its current strategy either performing uniform or proportional sampling with appropriate probabilities. This was an important simplification of the classical assumptions that up to now were used for proving convergence. However, the assumption of a bulletin board, implicitly makes use of global info for important characteristics of the system. Such info usually is unavailable on large scale networks as Internet. The main differences from [28] are the following. In [28] an infinite number jobs are assigned to an infinite number of machines, while their ratio remains constant. In [15] the resources are finite while the users are infinite. Also, in [28] agents exit from the system as soon as they allocate their jobs.

Contribution. Our motivation is to investigate the advantages and the limitations of a simple distributed protocol for congestion games on parallel links under very general assumptions on the latency functions. Hence we adopt a model of distributed computation allowing a limited amount of global knowledge, where in parallel every player can only select a link *uniformly at random* in each round and check its current latency. Migration decisions must be made concurrently on the basis only of the current latency of the resource (departure-link) to which a player is assigned and the current latency of the resource to which the player is about to move (destination-link). Thus, our replicator dynamics are based solely on local information. Migration decisions take advantage of *local coordination* amongst the players currently assigned to the *same* link, since at most one player is allowed to depart per link. The only global information about the latency functions is that they have a bounded slope. More precisely, we only assume that the latency functions satisfy the α-bounded jump condition (as in [7]).

Our notion of approximate pure Nash equilibrium, see Definition 2, is dictated by the very limited information that our model extracts, and is a bit different from similar approximate notions considered in previous work [7, 12] in an atomic setting, while it is close in nature to the stable state defined in [14, 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 links (even though this possibility is allowed by the definition of 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.

We present a simple oblivious protocol for this restricted model of distributed computation. According to our myopic protocol, 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 uses a simple local probabilistic rule that selects at most one (this is a local decision amongst users on the same link) potential migrant to defect from each link. 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 Rosental'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 result significantly extends the results in [4, 12] in the sense that (i) we consider arbitrary and unknown latency functions subject only to the α-bounded jump condition [7, 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 [22]. An interesting issue for further research is to extend its power by proportional sampling with respect to parameters that will favor its speed.

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, [21, 17] 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. [21] mostly considers congestion games on parallel links with identical users and convex delays. For this class of games, [21] 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. [17] 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 equiliria 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 [7] with the potential function of [17, 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 [7] 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 [13, 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

2.1 Model & target.

Model. There is a finite set of players $\{1, \ldots, n\}$ and a set of edges (or resources) $E = \{e_1, \ldots, e_m\}$. 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 $t = 0, \ldots, t^*$. It starts at round $t = 0$, where each player i selects myopically recourse $s_i(0) \in S_i$. In each subsequent round $t = 1, \ldots, t^*$, 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 2.2. That is, at round t the state $s(t) = \langle s_1(t), \ldots, s_n(t)\rangle \in S_1 \times \ldots \times S_n$ 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 $c_i(t) = \sum_{e \in s_i(t)} \ell_e(f_e(t))$ and let $\overline{c}(t) = \frac{1}{n}\sum_i c_i(t)$. Let the average delay of the resources be $\overline{\ell}(t) = \frac{1}{|E|}\sum_{e\in E} \ell_e(f_e(t)) = \frac{1}{m}\sum_{e \in E}\ell_e(f_e(t))$. Consider the value of Rosenthal's potential $\Phi(t) = \sum_{e\in E}\sum_{x=1}^{f_e(t)} \ell_e(x)$, and let $\overline{\Phi}(t) = \frac{\Phi(t)}{n}$. Clearly, if per round t only *one* player, say i, changes strategy $s_i(t)$ to $s_i(t+1)$, while the rest players remain fixed to their strategies: $\forall j \neq i, s_j(t) = s_j(t+1)$, then the profit $c_i(t) - c_i(t+1)$ of player i equals the potential drop $\Phi(t) - \Phi(t+1)$. This unilaterally strategy-changing sequential process will eventually converge at round t^* to a local minimum of $\Phi(\cdot)$.

Now, suppose that more that one players are allowed to migrate concurrently per round. Then, the process in not guaranteed to convergence towards to a Nash equilibrium. To see this, let player i sample an appealing strategy with *a priori* cost smaller than his cost $c_i(t)$ at hand. "A priori" means that the sampled cost is measured given that all other players *do not* change their strategies at hand. However, the actual or *a posteriori* cost $c_i(t+1)$ of i, i.e., when measured afterwards all concurrent player migrations have taken place, may turn out to be even higher than i's old cost $c_i(t)$. This may oscillate $\Phi(t+1)$ to a higher value than $\Phi(t)$. Such oscillations become more severe if the latency functions are arbitrary. We assume no latency-info other than the α-bounded jump condition:

Definition 1. *[7] Edge e satisfies the α-bounded jump condition if $\ell_e(f_e(t){+}1) \leq \alpha\ell_e(f_e(t)), \forall t \geq 0$.*

Here α is a positive constant, $\alpha > 1$, which contain a rich family of latencies, including polynomial ones. Lemma 1 is the main tool for handling probabilistically Rosenthal's potential and show fast convergence in expectation.

Our main result is that, despite its simplicity, `Greedy` in Section 2.2 does remarkably fast:

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

Theorem 1 follows easily (see Appendix A.1) from Theorem 2, see in turn its proof plan in Section 2.3. Here $\Phi_{\max}, (\Phi_{\min})$ denote the initial (final) value of the potential (value of the potential at an exact NE). Also $p = \Theta(1)$ is defined in Theorem 2. Our bicriterial equilibria (see [14, Def. 4]) follow.

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

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 link has latency greater than $\alpha\overline{\ell}(t)$ and no link with positive load has latency less than $\overline{\ell}(t)/\alpha$, 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 $\alpha\overline{\ell}(t)$) 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 $\overline{\ell}(t)/\alpha$). 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].

The idea on proving our main Theorem 1 is to show that each round which is not an almost-NE induces a potential drop of order of potential's current value and therefore at most logarithmical many rounds suffice. This is proved in our key Theorem 2.

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 bellow by a positive constant.*

The proof plan of this theorem is presented in Section 2.3. Its proof will be given in Section 2.7 which combines results proved in Section 2.4, 2.5 and 2.6.

2.2 Concurrent protocol `Greedy`

Initialization: Each player $i \in \{1, \ldots, n\}$ selects one random resource $e \in \{1, \ldots, m\}$.

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 $\vartheta = \Omega(1)$.

Note: For constants $\vartheta, \delta_\vartheta$ see Section 2.4 Lemma 2, Corollary 1 and also in Section 2.6 Case 1 and 2.

2.3 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 potential drop $\mathbb{E}[\Delta\Phi(t)]$ is a positive portion of the potential $\Phi(t)$ at hand. Since the minimum potential $\Phi_{\min}$ is a positive value, the total number of round is at most logarithmic in $\frac{\Phi_{\max}}{\Phi_{\min}}$. We present below how Sections 2.4, 2.5 and 2.6 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 2.4 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. Hence, showing large potential drop per round not in an almost-NE reduces to showing that $\sum_i \mathbb{E}[\Delta c_i(t)]$ is a positive number times $-\mathbb{E}[\Phi(t)]$.

This is achieved in Sections 2.5 and 2.6 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 2.5 shows that $\mathbb{E}[\Phi(t)]$ is a portion of $\mathbb{E}[\bar{\ell}(t)] \times m$. Having this, fast convergence reduces to showing that $\sum_i \mathbb{E}[\Delta c_i(t)]$ is a positive number times $-\mathbb{E}[\bar{\ell}(t)] \times m$ which is left to Section 2.6 & 2.7. At the end, Section 2.7 puts together Sections 2.4, 2.5 and 2.6 and completes the proof of our key Theorem 2.

2.4 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 2.2. 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`.

[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.

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

Proof. See Appendix A.2. $\square$

Lemma 2. *For every positive constant* δ*, if migration probability* ϑ *in step (3) of* `Greedy` *is* $\vartheta \leq \min\{\frac{\delta}{\alpha(\alpha-1)}, 1\}$*, 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 Appendix A.3 $\square$

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

Proof. See Appendix A.4 $\square$

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

It is easy to see this at round $t = 0$, since from the initialization of `Greedy` in Section 2.2 we get that the load of a bin is Binomially distributed:

$$\mathbb{E}[\overline{\ell}(0)] \leq \sum_{i=0}^{n}\binom{n}{i}\left(\frac{1}{m}\right)^i\left(1-\frac{1}{m}\right)^{n-i} \times \alpha^i \leq e^{\alpha\frac{n}{m-1}}e^{-\frac{n}{m}} = O(1), \text{ and}$$

$$\begin{aligned}
\mathbb{E}[\Phi(0)] &\leq \mathbb{E}\left[\sum_{e\in E}\sum_{x=1}^{f_e(0)}\ell_e(x)\right] \leq \sum_e \mathbb{E}[f_e(0)\ell_e(f_e(0))] \leq \mathbb{E}[f_e(0)\ell_e(f_e(0))] \times m\\
&\leq \left[\sum_{i=0}^{n}\binom{n}{i}\left(\frac{1}{m}\right)^i\left(1-\frac{1}{m}\right)^{n-i} \times i\alpha^i\right] \times m \leq \left[e^{\alpha\frac{n}{m-1}}e^{-\frac{n}{m}}\alpha\frac{n}{m-1+a}\right] \times m\\
&= \mathbb{E}[\overline{\ell}(0)]\alpha\frac{n}{m-1+a} \times m = O(\mathbb{E}[\overline{\ell}(0)] \times m) \qquad (1)
\end{aligned}$$

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 2.2). 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 super-martingale 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 Appendix A.5. $\square$

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 $\mathcal{B}_{L^*}(t) = E \setminus \mathcal{A}_{L^*}(t)$. The target of Lemma 4 is to show that $\mathcal{B}_{L^*}(t)$ is the subset of critical bins that will preserve similar bounds to (1) for each round $t > 0$ of `Greedy`.

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

Proof. Let $i_\ell = \min\{i : \alpha^i \geq \ell\}$. Recall $E_{\geq\nu}(t) = \{j \in E : f_j(t) \geq \nu\}$ defined in the proof of Lemma 3. Then:

$$\frac{1}{m}\sum_{e\in\mathcal{B}_{L^*}(t)} \mathbb{E}[\ell_e(f_e(t))] \leq \sum_i \Pr[f_e(t) = i \wedge \ell_e(f_e(t)) \geq L^*] \times \alpha^i$$
$$\leq \frac{|E_{\geq i_{L^*}}(t)|}{m} \times \sum_{i=i_{L^*}}^{n} \Pr[f_e(t) = i \,|\ell_e(f_e(t)) \geq L^*] \times \alpha^i \tag{2}$$

also

$$\frac{1}{m}\sum_{e\in\mathcal{B}_{L^*}(t)} \mathbb{E}[f_e(t)\ell_e(f_e(t))] \leq \sum_i \Pr[f_e(t) = i \wedge \ell_e(f_e(t)) \geq L^*] \times i\alpha^i$$
$$= \frac{|E_{\geq i_{L^*}}(t)|}{m} \times \sum_{i=i_{L^*}}^{n} \Pr[f_e(t) = i \,|\ell_e(f_e(t)) \geq L^*] \times i\alpha^i \tag{3}$$

Intuitively, we shall show that constant L^* is sufficiently high as to make edges in $\mathcal{B}_{L^*}(t)$ unlikely to receive so many balls as to deviate significantly from latency L^* per round $t > 0$ of `Greedy`. Also the remaining edges in $\mathcal{A}_{L^*}(t)$ are highly appealing, and balls accumulate in them per round, as long as their latency remains $< L^*$.

Under `Greedy`, we will show that Expressions (2) and (3) preserve similar bounds to the corresponding Expressions in (1), $\forall t > 0$. More precisely, $\forall t > 0$ `Greedy` draws 1 random ball per loaded bin (exactly as `Blind` does) and $n(t)$ totally. However, loaded critical bins with latency $\geq L^*$, are more unlikely, as it would be during `Blind`'s corresponding round, to receive random balls back. This is due to the super-martingale property in Corollary 2, that holds $\forall t > 0$ for all critical bins with latency $\geq L^*$. Intuitively, while the $n(t)$ drawn balls select random destinations, random destinations of latency $\geq L^*$ are "embargoed" by `Greedy`. Such a bias[5] towards random destinations of latency $< L^*$, induces Binomial tails for random destinations of latency $\geq L^*$.

From the above discussion we conclude that $\forall t$ it holds:

$$\frac{1}{m}\sum_{e\in\mathcal{B}_{L^*}(t)} \mathbb{E}[\ell_e(f_e(t))] \leq \mathbb{E}[\bar{\ell}(t)] = O(1) \tag{4}$$

also

$$\frac{1}{m}\sum_{e\in\mathcal{B}_{L^*}(t)} \mathbb{E}[f_e(t)\ell_e(f_e(t))] \leq \mathbb{E}[\bar{\ell}(t)] \times O(1) \tag{5}$$

□

[5] Observe that random destinations of latency $< L^*$ may become flooded by balls as the rounds evolute. However, as soon as any flooded destination e_0 becomes a critical one, its expected latency becomes $\leq L^*(\alpha + \delta_\theta) = O(1)$ by Lemma 2. And despite that e_0 has just become critical, the sparse way (sparser than `Blind`'s corresponding distribution of balls) that `Greedy` distributes balls on all critical destinations of latency $\geq L^*$ will keep e_0's expected latency constant in a way that similar expressions to the ones in (1) will remain true for all critical bins with latency $\geq L^*$, $\forall t > 0$.

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 Appendix A.8. □

2.6 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 $O(t)$ to $U(t)$ 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 2.4). 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 $U(t) = \{e \in E : \ell_e(f_e(t)) < (1-\delta)\bar{\ell}(t)\}$, while *overloaded* ones are $O(t) = \{e \in E : \ell_e(f_e(t)) \geq \alpha\bar{\ell}(t)\}$. Let us assume that we are not at an almost-NE because $|O(t)| \geq \varepsilon m$, with constant $\varepsilon \in (0, 1)$.

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

Proof. See Appendix A.9. □

Therefore, for every $e \in O(t)$, a player migrates from e to a link in $U(t)$ with probability at least $\theta\delta$ (see step (3) of `Greedy`, Section 2.2). Using Lemma 2 with $\theta = \varepsilon/4$, we obtain that the expected decrease in its cost is at least $\frac{\delta}{2}\alpha\bar{\ell}(t)$ (see Appendix A.6).

Given that k migrants switch from a link in $O(t)$ to a link in $U(t)$ we obtain that their expected cost-drop is at least $\frac{\delta}{2}\alpha\bar{\ell}(t)$ times their number k. Let $p_{O\to U}(k)$ the probability to have k such migrants. The expected number $\sum_k kp_{O\to U}(k)$ of such migrants is at least $\varepsilon\theta\delta m$, since for every $e \in O(t)$ with $|O(t)| \geq \epsilon m$, exactly 1 player migrates from e to a link in $U(t)$ with probability at least $\vartheta\delta$ (see Fact 4 and step (3) of `Greedy`, Section 2.2). Now, the unconditional on k expected cost-drop due to migrants switching from links in $O(t)$ to links in $U(t)$ is at least

$$\sum_k \left(\tfrac{\delta}{2}\alpha\bar{\ell}(t)k \times p_{O\to U}(k)\right) = \tfrac{\delta}{2}\alpha\bar{\ell}(t) \times \sum_k kp_{O\to U}(k) \geq \tfrac{\delta}{2}\alpha\bar{\ell}(t) \times \varepsilon\theta\delta m = \varepsilon\theta\tfrac{\delta^2}{2}\alpha m\bar{\ell}(t) \quad (6)$$

By (6) we finally prove (for Case 1) the result of this section:

$$\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq -\varepsilon\theta\tfrac{\delta^2}{2}\alpha \times \bar{\ell}(t)m \quad (7)$$

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

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

Proof. See Appendix A.10. $\square$

Since $|U(t)| \geq \varepsilon m$, a player migrates from each $e \in O(t)$ to a link in $U(t)$ with probability at least $\theta\varepsilon$ (see step (3) of `Greedy`, Section 2.2). Using Lemma 2 with $\theta = \frac{\varepsilon}{4\alpha}$, 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 Appendix A.7). Using Fact 5, we obtain that the expected cost-drop due to migrants leaving overloaded links $O(t)$ and entering $U(t)$ in round t is at least:

$$\theta\varepsilon \times \frac{\delta}{4} \sum_{e \in O(t)} \ell_e(f_e(t)) > \theta\varepsilon \times \frac{\delta}{4} \times \delta\, \overline{\ell}(t)m > \frac{\theta\varepsilon\delta^2}{4}\, \overline{\ell}(t)m \tag{8}$$

By (8) we finally prove (for Case2) the result of this section:

$$\sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq -\tfrac{\theta\varepsilon\delta^2}{4} \times m\overline{\ell}(t) \tag{9}$$

2.7 Proof of key Theorem 2.

Here we combine the results in Section 2.4, 2.5 and 2.6 and prove Theorem 2. From Section 2.4 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: (1) The *overloaded* links, with respect to the realization $\overline{\ell}(t)$, are $|O(t)| \geq \varepsilon m$. Then, we get from Expression (7) in Section 2.6 that $\mathbb{E}[\Delta\Phi(t)|\overline{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] < -\varepsilon\theta\frac{\delta^2}{2}\alpha \times \overline{\ell}(t)m$. (2) The *underloaded* links, with respect to the realization $\overline{\ell}(t)$, are $|U(t)| \geq \varepsilon m$. Then, we get from Expression (9) in Section 2.6 that $\mathbb{E}[\Delta\Phi(t)|\overline{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] < -\frac{\theta\varepsilon\delta^2}{4} \times \overline{\ell}(t)m$ In either Case 1 or 2 such that an almost-NE is not reached by realization $\overline{\ell}(t)$, we conclude from the above:

$$\mathbb{E}[\Delta\Phi(t)|\overline{\ell}(t)] \leq \sum_{i \in \mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] < -\tfrac{\theta\varepsilon\delta^2}{4} \times \overline{\ell}(t)m \tag{10}$$

Consider the space of all realizations $\overline{\ell}(t)$ not in an almost-NE due to $\geq \varepsilon m$ overloaded or underloaded links in round t. Let $p_{\overline{\ell}}(t)$ the probability to obtain a realization $\overline{\ell}(t)$ in this space. Removing the conditional on $\overline{\ell}(t)$, Expression (10) becomes:

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

From Fact 3 the above becomes: $\mathbb{E}[\Delta\Phi(t)] \leq -\frac{\theta\varepsilon\delta^2}{4} \times \frac{\mathbb{E}[\Phi(t)]}{r(1+y_t)+1+x_t}$, $r = n/m$ 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 $N = [n]$[6] partitioned into k coalitions $\{C_1, \ldots, C_k\}$, a set of resources $E = \{e_1, \ldots, e_m\}$, a strategy space $\Sigma_i \subseteq 2^E$ for each

[6] For every integer $n \geq 1$, $[n] \equiv \{1, \ldots, n\}$.

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 $\Sigma^{n_1} \times \cdots \times \Sigma^{n_k}$. 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 $s = (s_1, \ldots, s_k)$ 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 [13, 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[7]. 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 [17, Theorem 6] proves that Φ is an exact potential function for (even multi-commodity) congestion games with static coalitions and *linear* latencies. For sake of completeness, we include a proof sketch of [17, Theorem 6] in the Appendix, Section A.11. 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).

Theorem 6. *In a single-commodity linear congestion game with n players divided into k coalitions, the largest improvement ε-Nash dynamics starting from 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.*

Proof. See Appendix A.12. □

[7] For a vector $x = (x_1, \ldots, x_n)$, $x_{-i} \equiv (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$ and $(x_{-i}, x'_i) \equiv (x_1, \ldots, x_{i-1}, x'_i, x_{i+1}, \ldots, x_n)$.

References

1. Heiner Ackermann, Heiko Roeglin, and Berthold Voecking. On the impact of combinatorial structure on congestion games. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, 2006.
2. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05)*, pages 57–66, 2005.
3. B. Awerbuch, Y. Azar, A. Fiat, and T. Leighton. Making commitments in the face of uncertainty: How to pick a winner almost every time. In *Proceedings of the 28th ACM Symposium on the Theory of Computing, 1996, pp. 519–530.*, 1996.
4. Petra Berenbrink, Tom Friedetzky, Leslie Ann Goldberg, Paul Goldberg, Zengjian Hu, and Russell Martin. Distributed selfish load balancing. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 354–363, New York, NY, USA, 2006. ACM Press.
5. Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: on convergence to nash equilibria of regret-minimizing algorithms in routing games. In *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 45–52, New York, NY, USA, 2006. ACM Press.
6. J. Cao and C. Nyberg. An approximate analysis of load balancing using stale state information for servers in parallel. In *Proc. of the 2nd IASTED Int. Conf. on Communications, Internet, and Information Technology*, Nov 2003.
7. S. Chien and A. Sinclair. Convergece to Approximate Nash Equilibria in Congestion Games. In *Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*, page to appear, 2007.
8. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05)*, pages 67–73, 2005.
9. George Christodoulou, Vahab S. Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. In *STACS*, pages 349–360, 2006.
10. M. Dahlin. Interpreting Stale Load Information. *IEEE Transactions on Parallel and Distributed Systems*, 11(10), Oct 2001.
11. Eyal Even-Dar, Alex Kesselman, and Yishay Mansour. Convergence Time to Nash Equilibria. In *Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03)*, pages pages 502–513, 2003.
12. Eyal Even-Dar and Yishay Mansour. Fast convergence of selfish rerouting. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 772–781, 2005.
13. A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04)*, pages 604–612, 2004.
14. Simon Fischer, Harald Räcke, and Berthold Vöcking. Fast convergence to wardrop equilibria by adaptive sampling methods. In *STOC*, pages 653–662, 2006.
15. Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. In *PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, pages 276–283, New York, NY, USA, 2005. ACM Press.
16. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, 348:226–239, 2005.
17. D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic Congestion Games among Coalitions. In *Proc. of the 33th International Colloquium on Automata, Languages and Programming (ICALP '06)*, pages 573–584. Springer-Verlag, 2006.
18. Michel X. Goemans, Vahab S. Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *FOCS*, pages 142–154, 2005.
19. P. W. Goldberg. Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In *In Proc. of the twenty-third annual ACM symposium on Principles of distributed computing (PODC '04)*, pages 131–140, New York, NY, USA, 2004. ACM Press.
20. T. Grenager, R. Powers, and Y. Shoham. Dispersion games: general definitions and some specific learning results. In *AAAI*, 2002.
21. A. Hayrapetyan, É. Tardos, and T. Wexler. The Effect of Collusion in Congestion Games. In *Proc. of the 38th ACM Symp. on Theory of Computing (STOC '06)*, pages 89–98, 2006.
22. Samuel Ieong, Robert McGrew, Eugene Nudelman, Yoav Shoham, and Qixiang Sun. Fast and compact: A simple class of congestion games. In *AAAI*, pages 489–494, 2005.
23. A. Khanna and J. Zinky. The revised arpanet routing metric. In *SIGCOMM '89: Symposium proceedings on Communications architectures & protocols*, pages 45–56, New York, NY, USA, 1989. ACM Press.
24. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. In *Proc. of the 16thAnnual Symposium on Theoretical Aspects of Computer Science (STACS '99)*, pages 404–413. Springer-Verlag, 1999.
25. J. F. Kurose and K.W. Ross. *Computer Networking, a top down approch featuring the Internet*. Addison-Wesley Longman (3rd ed.), 2004.
26. Lavy Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.
27. Vahab S. Mirrokni and Adrian Vetta. Convergence issues in competitive games. In *APPROX-RANDOM*, pages 183–194, 2004.
28. Michael Mitzenmacher. How useful is old information? *IEEE Trans. Parallel Distrib. Syst.*, 11(1):6–20, 2000.
29. D. Monderer and L. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.

30. Ariel Orda, Raphael Rom, and Nahum Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.
31. J. Rexford. *Handbook of optimization in Telecommunications, chapter Route optimization in IP networks*. Kluwer Academic Publishers, 2005.
32. R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
33. J. W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.

A Appendix

A.1 Expected Time to Reach an Almost-NE

Assuming Theorem 2, we prove Theorem 1. Let $\Phi_{\max}$ denote the maximum value of the potential, and let $\Phi_{\min}$ denote the minimum value of the potential[8].

Notation. Let $\Phi(0)$ denote the potential of the initial state. For every round $t = 0, 1, \ldots$, we introduce a 0/1 random variable Z_t defined as:

$$Z_t = \begin{cases} 0 & \text{if GREEDY has not reached an almost-NE until round } t \\ 1 & \text{if GREEDY has reached an almost-NE no later than round } t \end{cases}$$

By definition, if $Z_\tau = 0$ for some round τ, then $Z_{\tau'} = 0$ for all rounds $\tau' \leq \tau$, and if $Z_\tau = 1$ for some round τ, then $Z_{\tau'} = 1$ for all rounds $\tau' \geq \tau$. We assume that $Z_0 = 0$, i.e. the initial state is not an almost-NE.

For a round $\tau \geq 1$, let $z_\tau \equiv \mathbb{P}\mathrm{r}[Z_\tau = 1 | Z_{\tau-1} = 0]$ denote the probability that GREEDY reaches an almost-NE in round τ given that it has not reached an almost-NE until round $\tau - 1$. Let $q_t \equiv \mathbb{P}\mathrm{r}[Z_t = 1]$ denote the probability that GREEDY has reached an almost-NE no later than round t. Applying Bayes' formula inductively, we obtain that $q_t = 1 - \prod_{\tau=1}^{t}(1 - z_\tau)$.

The Proof of Theorem 1. We start with a proposition bounding the conditional expectation of the potential at round t given that GREEDY has yet to reach an almost-NE.

Proposition 1. *For every round $t \geq 0$, $\mathbb{E}[\Phi(t)|Z_t = 0] \leq \frac{(1-p)^t}{1-q_t}\,\Phi(0)$.*

Proof. The proof uses induction on t. The proposition holds for $t = 0$ since $Z_0 = 0$, $q_0 = 0$, and $\mathbb{E}[\Phi(0)|Z_0 = 0] = \Phi(0)$.

We inductively assume that $\mathbb{E}[\Phi(t)|Z_t = 0] \leq \frac{(1-p)^t}{1-q_t}\,\Phi(0)$. Then Theorem 2 implies that

$$\mathbb{E}[\Phi(t+1)|Z_t = 0] \leq (1-p)\,\mathbb{E}[\Phi(t)|Z_t = 0] \leq \frac{(1-p)^{t+1}}{1-q_t}\,\Phi(0) \tag{11}$$

On the other hand,

$$\begin{aligned} \mathbb{E}[\Phi(t+1)|Z_t = 0] &\geq \mathbb{P}\mathrm{r}[Z_{t+1} = 0|Z_t = 0]\,\mathbb{E}[\Phi(t+1)|Z_{t+1} = 0] \\ &= (1 - z_{t+1})\,\mathbb{E}[\Phi(t+1)|Z_{t+1} = 0]\,, \end{aligned} \tag{12}$$

where we use that the potential is non-negative and that $Z_{t+1} = 0$ implies that $Z_t = 0$.

Combining (11) with (12), we obtain that

$$\mathbb{E}[\Phi(t+1)|Z_{t+1} = 0] \leq \frac{(1-p)^{t+1}}{(1-q_t)(1-z_{t+1})}\,\Phi(0)$$

Observing that $1 - q_{t+1} = (1 - q_t)(1 - z_{t+1})$ concludes the proof of the proposition. □

We proceed to bound from below the probability that GREEDY reaches an almost-NE in a sufficiently large number of rounds.

Proposition 2. GREEDY *reaches an almost-NE in $\lceil p^{-1}\ln(\beta^{-1}\Phi(0)/\Phi_{\min})\rceil$ rounds with probability at least $1 - \beta$.*

Proof. Let $t = \lceil p^{-1}\ln(\beta^{-1}\Phi(0)/\Phi_{\min})\rceil$. Assuming that $q_t < 1 - \beta$ and using Proposition 1, we obtain that

$$\mathbb{E}[\Phi(t)|Z_t = 0] < \tfrac{(1-p)^t}{\beta}\,\Phi(0) \leq \Phi_{\min}\,,$$

a contradiction. For the last inequality, we use that $1 - p \leq e^{-p}$ and that $t \geq p^{-1}\ln(\beta^{-1}\Phi(0)/\Phi_{\min})$. □

For $\beta = 1/2$, Proposition 2 implies that GREEDY reaches an almost-NE after an expected number of 2 phases, each consisting of $\lceil p^{-1}\ln(2\Phi_{\max}/\Phi_{\min})\rceil$ rounds.

[8] Notice that $\Phi_{\min} > 0$ because the α-bounded jump assumption is meaningful only if $\ell_e(1) > 0$ for all links e (otherwise, there is some link e with $\ell_e(k) = 0$ for all k and every player is bound to use e).

A.2 Proof of Lemma 1

It is helpful to construct the following directed graph $G(t) = (V(t), E(t))$ during round $t+1$. The vertices of $G(t)$ are the resources $V(t) = \{e_1, \dots, e_m\}$ and has $|\mathcal{A}(t)|$ directed edges. The directed edge $e_j \rightarrow e_k$ appears if a player moves from resource e_j to e_k during round $t+1$. According to `Greedy` each vertex has out-degree 1. That is, the edges $E(t)$ of $G(t)$ are the transactions made by players in $\mathcal{A}(t)$ per round. The *in(out)-degree* of a vertex is its number of in(out)coming edges, while its *degree* equals in-degree+out-degree. On each vertex $v \in V(t)$ with degree ≥ 1 we assign a color $\in$ {green, red, black} per round $t+1$ such that:

- Red are all vertices with in-degree 0 and out-degree 1. $\mathcal{A}_r(t)$ contains the players in $\mathcal{A}(t)$ that depart from a red vertex.
- Black are all vertices with in-degree ≥ 1 and out-degree 0.
- Green are all vertices with in-degree ≥ 1 out-degree 1. $\mathcal{A}_g(t)$ contains the players in $\mathcal{A}(t)$ that depart from a green vertex.

Observe that the contribution of terms in $\triangle[\Phi(t)]$ is only due to the colored vertices in $G(t)$. Specifically, red-vertices contribute to $\triangle[\Phi(t)]$ only negative terms, black-vertices contribute to $\triangle[\Phi(t)]$ only positive terms, green-vertices *may* contribute to $\triangle[\Phi(t)]$ only positive terms.

The negative terms in $\triangle[\Phi(t)]$ sum to:

$$\sum_{i \in \mathcal{A}(t)} (-c_i(t)) \tag{13}$$

To see this, first observe that each red-vertex e_j in $G(t)$ contributes to $\triangle[\Phi(t)]$ the negative term $-\ell_{e_j}(f_{e_j}(t)) = -c_i(t)$, where i is the player migrating from resource e_j during round $t+1$. Therefore, the transactions currently depicted in $G(t)$ *only* contribute to $\triangle[\Phi(t)]$ the following negative terms:

$$\sum_{i \in \mathcal{A}_r(t)} (-c_i(t)) = \sum_{i \in \mathcal{A}(t)} (-c_i(t)) - \sum_{i \in \mathcal{A}_g(t)} (-c_i(t)) \tag{14}$$

The crucial observation is that we can get our target (13) by plugging the missing terms $\sum_{i \in \mathcal{A}_g(t)}(-c_i(t))$ in (14) without affecting $\triangle[\Phi(t)]$ by the following trick: On each green resource we both add and subtract the corresponding term $-c_i(t)$ of the player $i \in \mathcal{A}_g(t)$ migrating from it, that is:

$$\triangle[\Phi(t)] = \triangle[\Phi(t)] + \sum_{i \in \mathcal{A}_g(t)} (-c_i(t)) + \sum_{i \in \mathcal{A}_g(t)} c_i(t) \tag{15}$$

We conclude that we have shown our target (13) without changing $\triangle[\Phi(t)]$.

The positive terms in $\triangle[\Phi(t)]$ sum to at most:

$$\sum_{i \in \mathcal{A}(t)} c_i(t+1) \tag{16}$$

This in turn can be shown as it follows. Each black vertex e_j with in-degree k contributes to $\triangle[\Phi(t)]$:

$$\sum_{x=1}^{k} \ell_{e_j}\left(f_{e_j}(t) + x\right) \leq k\ell_{e_j}\left(f_{e_j}(t) + k\right) = \sum_{i \in \mathcal{A}(t): s_i(t+1) = e_j} c_i(t+1) \tag{17}$$

Each green vertex e_j with in-degree k contributes to $\triangle[\Phi(t)]$:

$$\sum_{x=1}^{k-1} \ell_{e_j}\left(f_{e_j}(t) + x\right) \tag{18}$$

plus the corresponding term $\ell_{e_j}\left(f_{e_j}(t)\right)$ (added by trick) that appears in the rightmost summand in (15). Then (18) becomes:

$$\sum_{x=0}^{k-1} \ell_{e_j}\left(f_{e_j}(t) + x\right) \leq k\ell_{e_j}\left(f_{e_j}(t) + k - 1\right) = \sum_{i \in \mathcal{A}(t): s_i(t+1) = e_j} c_i(t+1) \tag{19}$$

Inequalities (17) and (19) show that (16) upper bounds the sum of positive terms in $\triangle[\Phi(t)]$ which proves our lemma when combined with (13).

A.3 Proof of Lemma 2

For every edge, GREEDY allows at most one player to migrate to a random link with probability ϑ. Hence, there are at most m candidate migrants and a link receives a migrant independently with probability ϑ/m. The distribution of the number of migrants in every link e is dominated by the binomial distribution $B(m, \vartheta/m)$.

Let e be an arbitrary destination link. Thus e receives some player, let it be player i. For every integer $k = 0, \ldots, m-1$, let Q_k denote the probability that the destination link e receives k additional players other than i. Since the number of candidate migrants (excluding player i) is $m-1$,

$$\begin{aligned} Q_k &\le \binom{m-1}{k}\left(\frac{\vartheta}{m}\right)^k\left(1-\frac{\vartheta}{m}\right)^{m-k} \\ &\le \frac{\vartheta^k}{k!}\left(1-\frac{1}{m}\right)^k\left(1-\frac{\vartheta}{m}\right)^{m-k} \\ &\le \frac{\vartheta^k}{k!}e^{-k/m}e^{-\vartheta(m-k)/m} \\ &\le \frac{\vartheta^k}{k!}e^{-\vartheta} \end{aligned} \tag{20}$$

The first inequality holds because the distribution of the number of additional migrants in e (other than player i) is dominated by the binomial distribution $B(m-1, \vartheta/m)$. For the second inequality, we use that $\binom{m-1}{k} \le \frac{(m-1)^k}{k!}$. For the third inequality, we use that $1 - x \le e^{-x}$ twice. For the last inequality, we use that $e^{-k(1-\vartheta)/m} \le 1$, since $\vartheta \le 1$.

Then the expected latency of the destination link e in the next round is bounded from above by:

$$\begin{aligned} \mathbb{E}[\ell_e(f_e(t+1))] &\le \sum_{k=0}^{\infty} Q_k\, \ell_e(f_e(t)+1+k) \\ &\le \ell_e(f_e(t)+1)\sum_{k=0}^{\infty} Q_k\, \alpha^k \\ &\le \ell_e(f_e(t)+1)\sum_{k=0}^{\infty}\frac{(\vartheta\alpha)^k}{k!}e^{-\vartheta} \\ &= e^{\vartheta(\alpha-1)}\ell_e(f_e(t)+1) \end{aligned}$$

The second inequality follows from the α-bounded jump assumption and the third inequality follows from (20). Using $\vartheta \le \frac{\delta}{\alpha(\alpha-1)}$, we obtain that $e^{\vartheta(\alpha-1)} \le 1 + \delta/\alpha$, which concludes the proof of the lemma.

A.4 Proof of Corollary 1

$\mathbb{E}[\Delta c_i(t)|\, c_i(t)] = \mathbb{E}[c_i(t+1)] - \ell_e(f_e(t))$. In Lemma 2 the migration probability ϑ is tuned such that $\forall i \in \mathcal{A}(t), \mathbb{E}[c_i(t+1)] \le \ell_{e'}(f_{e'}(t))(\alpha+\delta_\theta)$. By the selfish criterion of `Greedy` in Section 2.2 it holds $\ell_{e'}(f_{e'}(t))(\alpha+\delta_\theta) - \ell_e(f_e(t)) < 0$.

A.5 Proof of Lemma 3

Let $t \ge 0$ be any fixed round and let e be any fixed link with $\ell_e(f_e(t)) \ge \alpha^\nu$. We observe that

$$\mathbb{E}[f_e(t+1)] = f_e(t) + \mathbb{E}[\text{\#players coming in } e \text{ in } t] - \mathbb{E}[\text{\#players leaving } e \text{ in } t] \tag{21}$$

To establish the lemma, we show that if ν is sufficiently large, the expected number of players leaving e in t is no less than the expected number of players joining e in t.

Since $\ell_e(f_e(t)) \ge \alpha^\nu$, the α-bounded jump condition implies that e can receive players only from links in $E_{\ge\nu+1}(t) = \{j \in E : f_j(t) \ge \nu+1\}$. In particular, link e receives at most one player from each link in $E_{\ge\nu+1}(t)$ with probability ϑ/m. Therefore,

$$\mathbb{E}[\text{\#players coming in } e \text{ in } t] \le \vartheta|E_{\ge\nu+1}(t)|/m$$

On the other hand, since $\ell_e(f_e(t)) \ge \alpha^\nu$, the α-bounded jump condition implies that every link in $E_{\le\nu-2}(t) = \{j \in E : f_j(t) \le \nu-2\}$ satisfies the condition in step (3) of `Greedy` [9] Hence a player leaves e with probability at least $\vartheta|E_{\le\nu-2}(t)|/m$. Therefore,

$$\mathbb{E}[\text{\#players leaving } e \text{ in } t] \ge \vartheta|E_{\le\nu-2}(t)|/m$$

[9] For simplicity, we assume that the factor of $\alpha+\delta_\vartheta$ in step (3) of `Greedy` does not exceed α^2. In general, we have to use $E_{\le\nu-k-1}(t)$ (instead of $E_{\le\nu-2}(t)$) and $\nu \ge \lceil 2n/m\rceil + k$, where $k = \lceil\log_\alpha(1+\delta_\vartheta)\rceil$.

By (21), it suffices to show that for every integer $\nu \geq \lceil 2n/m \rceil + 1$, $|E_{\geq \nu+1}(t)| - |E_{\leq \nu-2}(t)| \leq 0$. Let $E_{\geq \nu-1}(t) = \{j \in E : f_j(t) \geq \nu - 1\}$. Then, $|E_{\geq \nu+1}(t)| \leq |E_{\geq \nu-1}(t)|$ and $|E_{\leq \nu-2}(t)| = m - |E_{\geq \nu-1}(t)|$. Moreover, $|E_{\geq \nu-1}(t)| \leq n/(\nu-1)$ by Markov's inequality. Therefore,

$$|E_{\geq \nu+1}(t)| - |E_{\leq \nu-2}(t)| \leq 2\,|E_{\geq \nu-1}(t)| - m \leq 0\,,$$

where we use that $|E_{\geq \nu-1}(t)| \leq m/2$ for all integers $\nu \geq \lceil 2n/m \rceil + 1$.

A.6 Cost decrease $\geq \frac{\delta}{2}\alpha\bar{\ell}(t)$

In Lemma 2 set $\theta = \frac{\varepsilon}{4}$ obtaining for the corresponding δ_θ that equals $\frac{\varepsilon}{4} = \frac{\delta_\theta}{\alpha(\alpha-1)} \Rightarrow \delta_\theta = \frac{\varepsilon\alpha(\alpha-1)}{4}$. The initial cost per user in $O(t)$ is $\geq \alpha\bar{\ell}(t)$. Each such user in $O(t)$ if it migrates to $U(t)$, by Corollary 1 its expected cost will be $\leq (\alpha+\delta_\theta)\,(1-\delta)\bar{\ell}(t)$, with $\delta = \frac{\varepsilon}{2}(\alpha-1)$, which is $\left(\alpha + \frac{\varepsilon\alpha(\alpha-1)}{4}\right)\left(1 - \frac{\varepsilon(\alpha-1)}{2}\right)\bar{\ell}(t) = \left(\alpha + \frac{\varepsilon\alpha(\alpha-1)}{4} - \frac{\varepsilon\alpha(\alpha-1)}{2} - \frac{\varepsilon^2\alpha(\alpha-1)^2}{8}\right)\bar{\ell}(t) < \alpha\left(1 - \frac{\varepsilon}{4}(\alpha-1)\right)\bar{\ell}(t) = \left(1 - \frac{\delta}{2}\right)\alpha\bar{\ell}(t)$, inducing a cost decrease $\geq \frac{\delta}{2}\alpha\bar{\ell}(t)$.

A.7 Cost decrease $\geq \frac{\delta}{4}\ell_e(f_e(t))$

In Lemma 2 set $\theta = \frac{\varepsilon}{4\alpha}$ obtaining for the corresponding δ_θ that equals $\frac{\varepsilon}{4\alpha} = \frac{\delta_\theta}{\alpha(\alpha-1)} \Rightarrow \delta_\theta = \frac{\varepsilon(\alpha-1)}{4}$. The initial cost per user on an edge $e_1 \in O(t)$ is $\geq (1+\delta)\,\bar{\ell}(t) = \left(1 + \frac{\varepsilon(\alpha-1)}{2\alpha}\right)\bar{\ell}(t)$, with $\delta = \frac{\varepsilon(\alpha-1)}{2\alpha}$. Therefore $\delta_\theta = \frac{\alpha}{2}\delta$. Each such user in $e_1 \in O(t)$ if it probes to $e_2 \in U(t)$, it holds:

$$\ell_{e_1}(t) \geq (1+\delta)\bar{\ell}(t) \geq \alpha(1+\delta)\ell_{e_2}(t)$$

where the 1st inequality holds from the definition of $O(t)$ and 2nd from $U(t)$. Therefore $\ell_{e_2}(t) \leq \frac{1}{\alpha(1+\delta)}\ell_{e_1}(t)$. Then, if user migrates to e_2 its expected latency will be $\leq \frac{\alpha+\delta_\theta}{\alpha(1+\delta)}\ell_{e_1}(t) = \frac{\alpha+\frac{\alpha}{2}\delta}{\alpha(1+\delta)}\ell_{e_1}(t) = \frac{1+\frac{\delta}{2}}{1+\delta}\ell_{e_1}(t)$. Its expected cost decrease will be $\left(1 - \frac{1+\frac{\delta}{2}}{1+\delta}\right)\ell_{e_1}(t) = \frac{\delta}{2(1+\delta)}\ell_{e_1}(t) \geq \frac{\delta}{4}\ell_{e_1}(t)$, since constant $\delta < 1$.

A.8 Proof of Fact 3

Let a sufficiently high constant L^* defined in Corollary 2 and $\mathcal{B}_{L^*}(t)$ the edges of latency $\geq L^*$ defined in Lemma 4. Observe that expected bound (5) for the overloaded edges in $\mathcal{B}_{L^*}(t)$ with latency $\geq L^*$ holds on every round of process `Greedy`, irrespectively if the round corresponds to an almost-NE or not. This expected bound will serve to the second summand in Expression (22) bellow as a bound to the contribution to the expected potential by the edges in $\mathcal{B}_{L^*}(t)$. Similarly, each underloaded edge in $\mathcal{A}_{L^*}(t) = E \setminus \mathcal{B}_{L^*}(t)$ encounters latency $< L^*$ per round. Since the number of users in edges in $\mathcal{A}_{L^*}(t)$ is $\leq n$ then the contribution to the potential by such edges is $< nL^*$ in the first summand in Expression (22) below, irrespectively if the round is an almost-NE. All in all, by linearity of expectation and the corresponding bounds nL^* and (5) we get:

$$\begin{aligned}
\mathbb{E}[\Phi(t)] &= \mathbb{E}\left[\sum_{e\in\mathcal{A}_{L^*}(t)}\sum_{i=1}^{f_e(t)}\ell_e(i) + \sum_{e\in\mathcal{B}_{L^*}(t)}\sum_{i=1}^{f_e(t)}\ell_e(i)\right] \leq \mathbb{E}\left[\sum_{e\in\mathcal{A}_{L^*}(t)}\sum_{i=1}^{f_e(t)}\ell_e(i) + \sum_{e\in\mathcal{B}_{L^*}(t)} f_e(t)\ell_e(f_e(t))\right]\\
&\leq nL^* + \sum_{e\in\mathcal{B}_{L^*}(t)}\mathbb{E}[f_e(t)\ell_e(f_e(t))] = \left(\frac{n}{m}L^* + \frac{1}{m}\sum_{e\in\mathcal{B}_{L^*}(t)}\mathbb{E}[f_e(t)\ell_e(f_e(t))]\right)m\\
&\leq \left(\frac{n}{m}L^* + \mathbb{E}[\bar{\ell}(t)]O(1)\right)m
\end{aligned} \tag{22}$$

Set $L^* = (1+y_t)\mathbb{E}[\bar{\ell}(t)]$ where $y_t = \Theta(1)$ since $L^* = O(1)$ and by Expression (4) it holds $\mathbb{E}[\bar{\ell}(t)] = O(1), \forall t$. Also, set $\mathbb{E}[\bar{\ell}(t)]O(1) = (1+x_t)\mathbb{E}[\bar{\ell}(t)]$, $x_t = \Theta(1)$. Finally, let $r = n/m = \Theta(1)$ and (22) becomes:

$$\mathbb{E}[\Phi(t)] \leq (r(1+y_t)+1+x_t)\,\mathbb{E}[\bar{\ell}(t)]m \Leftrightarrow \mathbb{E}[\bar{\ell}(t)]m \geq \frac{\mathbb{E}[\Phi(t)]}{r(1+y_t)+1+x_t} \tag{23}$$

A.9 Proof of Fact 4

For simplicity of notation, let $h = |O(t)|$, let $l = |U(t)|$, and let $n = m - h - l$ be the number of links with latency in $[(1-\delta)\bar{\ell}(t), \alpha\bar{\ell}(t)]$. Let us assume that $h \geq \varepsilon m$ and that $l < \delta m$. Then, $n > m - h - \delta m$. Calculating the average latency, we obtain a contradiction:

$$\begin{aligned} m\bar{\ell}(t) = \textstyle\sum_{e \in E} \ell_e(f_e(t)) &> h\,\alpha\bar{\ell}(t) + (m - h - \delta m)\,(1-\delta)\bar{\ell}(t) \\ &\geq \varepsilon m\,\alpha\bar{\ell}(t) + (1 - \varepsilon - \delta)m\,(1-\delta)\bar{\ell}(t) \\ &= m\bar{\ell}(t)(1 + \varepsilon(\alpha - 1) - \delta(1-\varepsilon) - \delta(1-\delta)) \geq m\bar{\ell}(t) \end{aligned}$$

The second inequality holds because $\alpha\bar{\ell}(t)\,(h - \varepsilon m) \geq (1-\delta)\bar{\ell}(t)\,(h - \varepsilon m)$, since $h \geq \varepsilon m$, $\alpha > 1$, and $\delta > 0$. The last inequality holds because for $\delta = \frac{\varepsilon}{2}(\alpha - 1)$, $\delta(1-\varepsilon) + \delta(1-\delta) \leq \varepsilon(\alpha - 1)$.

A.10 Proof of Fact 5

Working as in the proof of Fact 4, we obtain that

$$\sum_{e \in O(t)} \ell_e(f_e(t)) > (1 - \varepsilon/\alpha - (1-\varepsilon)(1+\delta))\,\bar{\ell}(t)m = \delta\,\bar{\ell}(t)m$$

A.11 Potential Function for Linear Congestion Games with Coalitions

Theorem 7. *([17]) Every linear congestion game with coalitions admits an exact potential function.*

Proof. We consider a game with a linear latency function $\ell_e(x) = a_e x + b_e$, $a_e, b_e \geq 0$, associated with each resource $e \in E$. Let $\{C_1, \dots, C_k\}$ be a set of coalitions, and let $s = (s_j)_{j \in [k]}$ be a pure strategies profile. The potential of s is $\Phi(s) = \frac{1}{2}[C(s) + W(s)]$, where

$$C(s) = \sum_{e \in E} f_e(s)\ell_e(f_e(s)) = \sum_{e \in E} (a_e f_e^2(s) + b_e f_e(s))\,, \text{ and}$$

$$W(s) = \sum_{e \in E} \sum_{j=1}^{k} f_e(s_j)\ell_e(f_e(s_j)) = \sum_{e \in E} \sum_{j=1}^{k} (a_e f_e^2(s_j) + b_e f_e(s_j))$$

Let j be a coalition changing its strategy from s_j to s'_j, and let $s' = (s_{-j}, s'_j)$ be resulting pure strategies profile. The difference in the total delay of coalition j is[10]:

$$\tau_j(s') - \tau_j(s) = \sum_{e \in E} (f_e(s'_j) - f_e(s_j))[a_e(f_e(s_{-j}) + f_e(s'_j) + f_e(s_j)) + b_e] \tag{24}$$

The difference in the potential function components is:

$$C(s') - C(s) = \sum_{e \in E} (f_e(s'_j) - f_e(s_j))[a_e(2\,f_e(s_{-j}) + f_e(s'_j) + f_e(s_j)) + b_e] \tag{25}$$

and

$$W(s') - W(s) = \sum_{e \in E} (f_e(s'_j) - f_e(s_j))[a_e(f_e(s'_j) + f_e(s_j)) + b_e] \tag{26}$$

Combing (25), (26), and (24), we obtain that $\Phi(s') - \Phi(s) = \tau_j(s') - \tau_j(s)$. Therefore, Φ is an exact potential for linear congestion games with coalitions. ◻

[10] In the following, we repeatedly use that for every pure strategies profile s, coalition j, and resource e, $f_e(s) = f_e(s_{-j}) + f_e(s_j)$.

A.12 Proof of Theorem 6

The outline of the proof is similar to that of [7, Theorem 3.4], which holds for *symmetric* congestion games only. However, coalitions may be of different size, in which case the game is *asymmetric*. Hence, we have to extend the technique of [7] and bound the effect of coalitions of different size. On the other hand, our result holds for a more restricted class of latency functions compared to that in [7].

Let $\{C_1, \dots, C_k\}$ be a set of coalitions, and let $s = (s_j)_{j\in[k]}$ be a pure strategies profile that is not an ε-Nash equilibrium. We prove that every ε-move dictated by the largest improvement dynamics decreases the potential by at least $\frac{\varepsilon(1-\varepsilon)}{kr(r+1)}\Phi(s)$. This implies the theorem, since the potential is initially $\Phi(s_0)$ and Φ is a non-negative integral function.

Since $\Phi(s) \le \sum_{j=1}^{k} \tau_j(s)$, there is some coalition of total delay at least $\Phi(s)/k$. Let C_i be a coalition of maximum total delay in s. Clearly, $\tau_i(s) \ge \Phi(s)/k$. Let s'_i be C_i's best response to s_{-i}. We distinguish between two cases depending on whether $(1-\varepsilon)\tau_i(s) > \tau_i(s_{-i}, s'_i)$, i.e. C_i has an ε-move available, or not.

If C_i has an ε-move available, the next move decreases the potential by at least $\varepsilon\Phi(s)/k$. More precisely, if C_i moves, then

$$\Phi(s) - \Phi(s_{-i}, s'_i) = \tau_i(s) - \tau_i(s_{-i}, s'_i) > \varepsilon\tau_i(s) \ge \varepsilon\,\Phi(s)/k$$

The equality holds because Φ is an exact potential (see also the proof of Theorem 7). The first inequality follows from the hypothesis that C_i makes an ε-move. The last inequality follows from the definition of C_i as a maximum cost coalition in s. If instead of C_i, some other coalition C_j moves from s_j to s'_j, by the definition of the largest improvement dynamics, $\tau_j(s) - \tau_j(s_{-j}, s'_j) \ge \tau_i(s) - \tau_i(s_{-i}, s'_i)$, and the potential decreases by at least $\varepsilon\,\Phi(s)/k$.

If C_i does not have an ε-move available, let C_j be the coalition that moves from s_j to s'_j and hence decreases the potential by $\varepsilon\tau_j(s)$. We show that $\tau_j(s) \ge \frac{(1-\varepsilon)}{kr(r+1)}\Phi(s)$. Therefore, the potential decreases by at least $\frac{\varepsilon(1-\varepsilon)}{kr(r+1)}\Phi(s)$.

Let $\tilde{s}_j$ be the strategy of coalition C_i obtained by taking $\lceil n_i/n_j \rceil$ copies of s_j. More precisely, $\tilde{s}_j$ is obtained by assigning at most $\lceil n_i/n_j \rceil$ players from C_i to each strategy s_j^ν, $\nu \in C_j$, until all players in C_i are assigned to some strategy in Σ. Then,

$$\begin{aligned}
\tau_i(s_{-i}, \tilde{s}_j) &\le \sum_{e\in E} f_e(\tilde{s}_j)\,\ell_e(f_e(s) + f_e(\tilde{s}_j)) \\
&\le \sum_{e\in E} \lceil n_i/n_j \rceil\, f_e(s_j)\,\ell_e(f_e(s_{-j}) + (\lceil n_i/n_j \rceil + 1) f_e(s_j)) \\
&\le \lceil n_i/n_j \rceil (\lceil n_i/n_j \rceil + 1) \sum_{e\in E} f_e(s_j)\,\ell_e(f_e(s_{-j}) + f_e(s_j)) \\
&\le r(r+1)\tau_j(s)
\end{aligned}$$

The second inequality holds because by the definition of $\tilde{s}_j$, $f_e(\tilde{s}_j) \le \lceil n_i/n_j \rceil\, f_e(s_j)$ for every resource e. The third inequality follows from the linearity of the latency functions. The last inequality holds because $\lceil n_i/n_j \rceil \le r$.

Therefore, $\tau_j(s) \ge \frac{\tau_i(s_{-i},\tilde{s}_j)}{r(r+1)}$. Since C_i does not have an ε-move available, $(1-\varepsilon)\tau_i(s) \le \tau_i(s_{-i}, \tilde{s}_j)$, which implies that $\tau_i(s_{-i}, \tilde{s}_j) \ge (1-\varepsilon)\Phi(s)/k$ and that $\tau_j(s) \ge \frac{1-\varepsilon}{kr(r+1)}\Phi(s)$. Thus, as soon as C_j switches from s_j to s'_j, the potential decreases by at least $\frac{\varepsilon(1-\varepsilon)}{kr(r+1)}\Phi(s)$.