Atomic congestion games: fast, myopic *and* concurrent

(Invited paper to a special issue of TOCS dedicated to the SAGT'08 best papers)

D. Fotakis

Dept. of Information and Communication Systems Engineering, U. Aegean, fotakis@aegean.gr

A.C. Kaporis

Dept. of Computer Engineering and Informatics, U. Patras, kaporis@ceid.upatras.gr

P.G. Spirakis

Dept. of Computer Engineering and Informatics, U. Patras

& RA Computer Technology Institute, 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 [37] showed that pure Nash equilibria on atomic congestion games correspond to local minima of a natural potential function. Twenty years later, Monderer and Shapley [32] proved that congestion games are equivalent to potential games. Many recent contributions have provided considerable insight into the structure and efficiency (e.g. [18, 2, 9, 21]) and tractability [15, 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 [37] 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 equilibrium. However, this may require an exponential number of steps, since computing a pure Nash equilibrium of a congestion game is *PLS-complete* [15].

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 [15]. 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 [25, 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. [8] considers *symmetric* congestion games with a weak restriction on latency functions (bounded-jump latencies) 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 [10, 14, 21, 22, 29, 30, 35]. 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 [15] 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. 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 [40] perspective, which studies conditions under which a population of agents may (or may not) reach stable states, see also [38] 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. 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 [26, 28, 36].

Let us first focus to the discrete concurrent setting. The work in [13] considers n players concurrently sampling 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 [27]. This migration protocol, although concurrent, is not completely decentralized, since it uses global information 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 [23] was prior to [13], were a series of similar concurrent protocols were validated.

In [5] 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 [13] were related, but the important aspect of the analysis in [5] 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. The work in [39] 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 important papers [6, 16] provide strong intuition on this subject. More precisely [16] 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 [16] 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 [16] studies specific replication policies designated to yield fast convergence, in [6] 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 [6]: "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 [31, 17] 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 [7, 12, 3]. More precisely, in [17] 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 [31] are the following. In [31] an infinite number jobs are assigned to an infinite number of machines, while their ratio remains constant. In [17] the resources are finite while the users are infinite. Also, in [31] 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 resources 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 resource *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-resource) to which a player is assigned and the current latency of the resource to which the player is about to move (destination-resource). 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* resource, since at most one player is allowed to depart per resource. 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 latency jump bound (see also [8]).

Our notion of ε-approximate equilibrium ((ε, α)-EQ), in Def. 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 [8, 13] in an atomic setting, while it is close in nature to the stable state defined in [16, Def. 4] for the Wardrop model. An (ε, α)-EQ is a state where at most εm resources 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) equilibria for our fully myopic model of migration described above. In particular, an (ε, α)-EQ guarantees that unless a player uses an overloaded resource (i.e. a resource with latency considerably larger than the average latency), the probability that she finds (by uniform sampling) a resource to migrate and significantly improve her latency is at most $o(1)$. Furthermore, it is unlikely that any (ε, α)-EQ reached by our protocol assigns any number of players to overloaded resources (even though this possibility is allowed by the definition of an (ε, α)-EQ). As it will become clear from the analysis in Section 2.6, the reason that users do not accumulate on overloaded resources, is that the number of players on such resources is a strong super-martingale. Initially, with high probability, each resource has load $O(\log n)$. Then we get that in expectedly $O(\log n)$ rounds the overloaded resources will drain from users.

We present (Sect. 2.3) a simple oblivious protocol for this restricted model of distributed computation. According to our myopic protocol, in parallel each player selects a resource uniformly at random in each round and checks whether she can significantly decrease her latency by moving to the chosen resource. 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 resource) potential migrant to defect from each resource. We prove (Th. 1) that if the number of players is $\Theta(m)$, the protocol reaches an almost-NE in $O(\log(\mathbb{E}[\Phi(0)]/\Phi_{\min}))$ time, where $\mathbb{E}[\Phi(0)]$ is Rosental's expected potential value as the game starts and $\Phi_{\min}$ is the corresponding value at a NE. The proof of convergence (Sec. 2.6) is technically involved and interesting and comprises the main technical merit of this work.

Our result significantly extends the results in [5, 13] in the sense that (i) we consider arbitrary and unknown latency functions subject only to the α latency jump bound [8, Section 2], (ii) it requires no other global information. 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 [25]. 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, [24, 19] 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. [24] mostly considers congestion games on parallel links with identical users and convex latency functions. For this class of games, [24] 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. [19] presents a potential function for linear congestion games with coalitions.

Contribution. In this setting, we present (Sec. 3) an upper bound (Th. 5) on the rate of (sequential) convergence to approximate NE 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 NE. We consider sequences of ε-moves, i.e. selfish deviations that improve the coalitions' total delay by a factor greater than ε. Combining the approach of [8] with the potential function of [19, Theorem 6], we show that if the coalition with the largest improvement in its total delay moves in every round, an approximate NE is reached in a small number of steps.

More precisely, we prove (Th. 5) 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 results of [8] 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 *non-symmetric* congestion games.

This bound implies that in network congestion games, where a coalition's best response can be computed in polynomial time by a min-cost flow computation [15, Theorem 2], an approximate 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. when $k = \Theta(1)$ and $r = \Theta(1)$, the number of ε-moves to reach an approximate NE is logarithmic in the potential of the initial state.

2 Concurrent atomic congestion games

2.1 Model

There is a finite set of n players $N = \{1, \ldots, n\}$ and a set of m edges (or resources) $E = \{e_1, \ldots, e_m\}$, where $n = O(m)$ and $r = n/m = O(1)$. The strategy space S_i of player i is E. The game consists of a sequence of rounds $t = 0, \ldots, T^*$, T^* is the round they reach an (ε, α)-EQ as Def. 2 for the first time. The strategy of player $i \in N$ at round t is $s_i(t) \subseteq E$. We study *singleton* games, i.e., $|s_i(t)| = 1, \forall i \in N$. 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 n-tuple 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 $\ell_e(f_e(t)) \geq 0$ measuring the common delay of players on e at round t, increasing on load $f_e(t)$. 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 latency* of the resources be $\overline{\ell}(t) = \frac{1}{m} \sum_{e \in E} \ell_e(f_e(t))$. Consider the value of Rosenthal's potential [37] $\Phi(t) = \sum_{e \in E} \sum_{x=1}^{f_e(t)} \ell_e(x)$. We assume no latency-info other than the *α-latency jump bound*:

Definition 1 *[8] Let* $\alpha = \min\{a | \; \forall x = 0, \ldots, n, \forall e \in E$ *it holds* $\ell_e(x+1) \leq a \ell_e(x)\}$. *Then each resource* $e \in E$ *satisfies the* α-latency jump bound.

This condition imposes a minor restriction on the increase-rate of the latency function $\ell_e()$ 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 definition of bicriteria equilibrium, reminiscent to [16, Def. 4], follows.

2.2 Bicriteria equilibrium

Definition 2 *An* (ε, α)-EQUILIBRIUM ((ε, α)-EQ) *is a state where* $\leq \varepsilon m$ *loaded resources have latency* $> \alpha\overline{\ell}(t)$ *and* $\leq \varepsilon m$ *loaded resources have latency* $< \frac{1}{\alpha}\overline{\ell}(t)$*, where* α *is the latency jump bound.*

Taking into account the very limited info that our protocol extracts per round, our analysis suggests that an (ε, α)-EQ is a meaningful notion of a stable state that can be reached quickly. In particular, the (ε, α)-EQ 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) resources to migrate and significantly improve their cost is small.

More precisely, in an exact NE, no loaded resource has latency greater than $\alpha\overline{\ell}(t)$ and no resource with positive load has latency less than $\overline{\ell}(t)/\alpha$, while the definition of an (ε, α)-EQ imposes the same requirements on all but εm resources. Hence the notion of an (ε, α)-EQ is a relaxation of the notion of an exact NE. In addition, a player not assigned to an overloaded resource (i.e. a resource 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 resource (i.e. a resource with latency less than $\overline{\ell}(t)/\alpha$). Therefore, in an (ε, α)-EQ, the probability that a player not assigned to an overloaded resource samples a resource where she can migrate and significantly decrease her cost is $\leq \varepsilon m$. Furthermore, it is unlikely that the (ε, α)-EQreached by our protocol assigns a large number of players to overloaded resources [1].

2.3 Protocol `Greedy`

At $t = 0$, concurrently each player i selects a random strategy (resource) $s_i(0) \in S_i$, while for $t > 0$, each player updates her $s_i(t)$ to $s_i(t+1)$ according to protocol `Greedy` illustrated below (recall α in Def. 1):

During each round $t \geq 1$, do in parallel $\forall e \in E$:

1. Select 1 player i from resource e at random, with latency $\ell_e(f_e(t))$.
2. Player i samples for a destination resource e' u.a.r. over E, with latency $\ell_{e'}(f_{e'}(t))$.
3. If $\ell_{e'}(f_{e'}(t))(\alpha + \delta) < \ell_e(f_e(t))$, then move player i from e to e' with probability ϑ.

[1] Due to the initial random allocation of the players to the resources, the overloaded resources (if any) receive $O(\log n)$ players with high probability. Lemma 4 shows that the number of players on any overloaded resource is a strong super-martingale during each round. Thus, such overloaded resources will drain from users in expectedly $O(\log n)$ rounds.

We show that Greedy quickly (Th. 1) reaches an (ε, α)-EQ (Def. 2), where ϑ, δ are parameters tuned in Theorem 1.

2.4 Greedy: insights from distributed computing and load balancing

In our work, there are quite a few points where our research draws from advances in other fields of computing, beyond that of algorithmic game theory.

A key such point is the nature of the protocol that decides who migrates between resources and how, as well as the extent to which such migrations effectively and efficiently achieve some notion of optimality. The field that has been most influential in that respect is that of load balancing, where key results [11] suggest that migration protocols are realistic when they assume that (now, we switch to the game nomenclature) a number of players moving from one resource at a given time point (round) actually move to the *same* target, and are not distributed amongst more than one target. This differentiation is described as the contrast between *diffusion* and *dimension exchange* methods, where the latter impose that a resource will only communicate (sample) with *one* potential target resource, to determine where to allow some of its migrants to move to (if at all). It is important to note that this assumption improves the robustness of the migration protocol since, when considering which players to move out of a resource, we do not need to collect expensive information (as is the case, for example, in [13]) from *all* available resources but we just focus on sampling one potential target. To appreciate this robustness, consider what would happen in a network where we might need to sample many resources, yet find that many of the links seem to be broken, as is quite likely of course.

The justification for our protocols can be further seen in [20], where it is argued why a resource cannot be expected to communicate in parallel with other resources, leading to the observation that sequential communication means that all migrants from a source will all go to the same target. Morever, also according to [20], we note that our protocol indeed realistically assume that only local information is made available to the migrating candidates; note that, in stark contrast to this recommendation, [13] assume that players have access to accurate global statistics (like average load) to compute their next move.

A further justification for our protocol is the design pattern discussion in [4], where analogues are drawn to several biological processes that have influenced the design of distributed computing protocols and algorithms, and where a central recurring theme is the identification of processes that rely on strictly local information yet manage to achieve some notion of effective global behavior.

2.5 Main results

Our main result is Theorem 1, establishing `Greedy`'s convergence to an (ε, α)-EQ (Def. 2).

Theorem 1 `Greedy` *reaches an* (ε, α)-EQ *in expectedly* $O\left(\frac{4r}{\vartheta\varepsilon\delta^2}\ln(\frac{2\mathbb{E}[\Phi(0)]}{\Phi_{\min}})\right)$ *rounds, with* $\vartheta = \frac{\varepsilon}{4\alpha}$, $\delta = \frac{\varepsilon(\alpha-1)}{2\alpha}$, $\Phi_{\min}$ *the optimal potential,* $r = n/m$, $\mathbb{E}[\Phi(0)]$ *the average potential at round* $t = 0$*, and* $\alpha > 1$ *the latency jump bound (Def. 1).*

Theorem 1 follows by inductive application of Theorem 2.

Theorem 2 $\mathbb{E}[\Delta\Phi(t)] \leq -\frac{\vartheta\varepsilon\delta^2}{4r} \times \Theta(1) \times \mathbb{E}[\Phi(t)]$*, for each round* t *not in an* (ε, α)-EQ*, with* $\vartheta = \frac{\varepsilon}{4\alpha}$, $\delta = \frac{\varepsilon(\alpha-1)}{2\alpha}$, $\Phi_{\min}$ *the optimal potential,* $r = n/m$, $\mathbb{E}[\Phi(0)]$ *the average potential at round* $t = 0$*, and* $\alpha > 1$ *the latency jump bound (Def. 1).*

The proof is in Section 2.6.4. The line of thought is presented in Section 2.6.

2.6 Proof of convergence of `Greedy` - Overview

The idea behind main Theorem 1 is to show that, starting from $\mathbb{E}[\Phi(0)]$, per round t of `Greedy` not in an (ε, α)-EQ, the expected $\mathbb{E}[\Delta\Phi(t)]$ potential drop 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 $\frac{\mathbb{E}[\Phi(t)]}{\Phi_{\min}}$. We present below how Sections 2.6.1, 2.6.2 and 2.6.3 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 (ε, α)-EQ, and prove key Theorem 2. Section 2.6.1 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* [33, Def. 4.7]. Hence, showing large potential drop per round not in an (ε, α)-EQ 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 2.6.2 and 2.6.3 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 2.6.2 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 2.6.3 & 2.6.4. At the end, Section 2.6.4 puts together Sections 2.6.1, 2.6.2 and 2.6.3 and completes the proof of our key Theorem 2.

2.6.1 Showing that $\mathbb{E}[\Delta\Phi(t)] \leq \sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq 0$

Let $\mathcal{A}(t)$ the migrants allowed in step (3) of `Greedy` in Section 2.3.

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

Proof. 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 $V(t) = \{e_1, \ldots, e_m\}$ and $G(t)$ has $|\mathcal{A}(t)|$ directed edges. The directed edge $e_j \to e_k$ appears if a player moves from resource e_j to e_k during round $t+1$. The *in(out)-degree* of a vertex is its number of in(out)coming edges, while its *degree* equals in-degree+out-degree. That is, the edges $E(t)$ of $G(t)$ are the transactions made by players in $\mathcal{A}(t)$ per round. According to `Greedy` each vertex has out-degree 1. 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 and 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:

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

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:

$$\sum_{i\in\mathcal{A}_r(t)} (-c_i(t)) = \sum_{i\in\mathcal{A}(t)} (-c_i(t)) - \sum_{i\in\mathcal{A}_g(t)} (-c_i(t)) \tag{2}$$

The crucial observation is that we can get our target (1) by plugging the missing terms $\sum_{i\in\mathcal{A}_g(t)}(-c_i(t))$ in (2) 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:

$$\triangle[\Phi(t)] = \triangle[\Phi(t)] + \sum_{i\in\mathcal{A}_g(t)} (-c_i(t)) + \sum_{i\in\mathcal{A}_g(t)} c_i(t) \tag{3}$$

We conclude that we have shown our target (1) without changing $\triangle[\Phi(t)]$.

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

$$\sum_{i\in\mathcal{A}(t)} c_i(t+1) \tag{4}$$

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

$$\sum_{x=1}^{k} \ell_{e_j}\left(f_{e_j}(t)+x\right) \leq k\ell_{e_j}\left(f_{e_j}(t)+k\right) = \sum_{i\in\mathcal{A}(t):s_i(t+1)=e_j} c_i(t+1) \tag{5}$$

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

$$\sum_{x=1}^{k-1} \ell_{e_j}\left(f_{e_j}(t)+x\right) \tag{6}$$

plus the corresponding term $\ell_{e_j}\left(f_{e_j}(t)\right)$ (added by trick) that appears in the rightmost summand in (3). Then (6) becomes:

$$\sum_{x=0}^{k-1} \ell_{e_j}\left(f_{e_j}(t)+x\right) \leq k\ell_{e_j}\left(f_{e_j}(t)+k-1\right) = \sum_{i\in\mathcal{A}(t):s_i(t+1)=e_j} c_i(t+1) \tag{7}$$

Inequalities (5) and (7) show that (4) upper bounds the sum of positive terms in $\triangle[\Phi(t)]$ which proves our lemma when combined with (1). □

Lemma 1 and linearity of expectation yield $\mathbb{E}[\Delta\Phi(t)] \leq \sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)]$. It remains to show

$$\sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] < 0$$

This follows by Lemma 2 & Corollary 1.

Lemma 2 *If the migration probability*[2] *of* `Greedy` *is* $\vartheta \leq \min\{\frac{\delta}{\alpha(\alpha-1)}, 1\}$ *then*

$$\mathbb{E}[\ell_e(f_e(t+1))] \leq (1+\delta/\alpha)\ell_e(f_e(t)+1) \leq (\alpha+\delta)\ell_e(f_e(t)), \forall\delta > 0$$

Proof. For every resource, `Greedy` allows at most one player to migrate to a random resource with probability ϑ. Hence, there are at most m candidate migrants and a resource receives a migrant independently with probability $\leq \vartheta/m$. The distribution of the number of migrants in

[2] In Sec. 2.3 `Greedy` reaches an (ε,α)-EQ with migration probability $\vartheta = \frac{\varepsilon}{4\alpha}$ and $\delta = \frac{\varepsilon(\alpha-1)}{2}$, where $\alpha > 1$ is the latency jump bound (Def. 1).

every resource e is dominated by the binomial distribution $B(m, \vartheta/m)$. Let e be an arbitrary destination resource. 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 resource e receives k additional players other than i. Since the number of candidate migrants (excluding player i) is $m-1$,

$$\begin{aligned} Q_k &\leq \binom{m-1}{k}\left(\frac{\vartheta}{m}\right)^k\left(1-\frac{\vartheta}{m}\right)^{m-k} \\ &\leq \frac{\vartheta^k}{k!}\left(1-\frac{1}{m}\right)^k\left(1-\frac{\vartheta}{m}\right)^{m-k} \\ &\leq \frac{\vartheta^k}{k!}e^{-k/m}e^{-\vartheta(m-k)/m} \\ &\leq \frac{\vartheta^k}{k!}e^{-\vartheta} \end{aligned} \tag{8}$$

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 $\binom{m-1}{k} \leq \frac{(m-1)^k}{k!}$. For the third inequality, we use that $1-x \leq e^{-x}$ twice. For the last inequality, we use that $e^{-k(1-\vartheta)/m} \leq 1$, since $\vartheta \leq 1$. Then the expected latency of the destination resource e in the next round is bounded from above by:

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

The second inequality follows from the α latency jump bound (Def. 1) and the third inequality follows from (8). Using $\vartheta \leq \frac{\delta}{\alpha(\alpha-1)}$, we obtain that $e^{\vartheta(\alpha-1)} \leq 1+\delta/\alpha$, which concludes the proof of the lemma. □

Corollary 1 *If* `Greedy` *'s migration probability is* $\vartheta = \frac{\varepsilon}{4\alpha}$ *then*

$$\mathbb{E}[\Delta c_i(t)|\, c_i(t)] \leq \ell_{e'}(f_{e'}(t))(\alpha+\delta_\vartheta) - c_i(t) \leq 0, \textit{ with } \delta_\vartheta = \frac{\varepsilon(\alpha-1)}{2}$$

$\forall$ *player* $i \in \mathcal{A}(t)$ *that moves from* e *to* e'.

Proof. $\forall$ player $i \in \mathcal{A}(t)$ that moves from e to e', her current cost is $c_i(t) = \ell_e(f_e(t))$. Then $\mathbb{E}[\Delta c_i(t)|\ c_i(t)] = \mathbb{E}[c_i(t+1)] - \ell_e(f_e(t))$. The migration probability $\vartheta = \frac{\varepsilon}{4\alpha}$ and Lemma 2 yields $\forall i \in \mathcal{A}(t), \mathbb{E}[c_i(t+1)] \leq \ell_{e'}(f_{e'}(t))(\alpha + \delta_\vartheta)$, with $\delta_\vartheta = \frac{\varepsilon(a-1)}{2}$. In Section 2.3, `Greedy` moves a player only if it holds $\ell_{e'}(f_{e'}(t))(\alpha + \delta_\vartheta) - \ell_e(f_e(t)) < 0$. □

2.6.2 Showing that $\mathbb{E}[\Phi(t)] \leq \Theta(1) \times \mathbb{E}[\bar{\ell}(t)] \times m$

The lemma bellow is a consequence of the strong tail bounds on the load of any resource at round $t = 0$. Since, the initialization of `Greedy` in Section 2.3 implies that the load of any resource is Binomially distributed.

Lemma 3 $\mathbb{E}[\Phi(0)] = \Theta(1) \times \mathbb{E}[\bar{\ell}(0)]m$

Proof. At round $t = 0$, the load $f_e(0)$ of a resource e is Binomial random variable. Thus:

$$\mathbb{E}[\Phi(0)] = \mathbb{E}\left[\sum_{e \in E}\sum_{x=1}^{f_e(0)} \ell_e(x)\right] \leq \sum_{e \in E} \mathbb{E}[f_e(0)\ell_e(f_e(0))] = \sum_{e \in E}\sum_{i=0}^{n} \Pr[i] i \ell_e(i)$$
$$\text{with } \Pr[i] = \binom{n}{i}\left(\frac{1}{m}\right)^i \left(1 - \frac{1}{m}\right)^{n-i} \tag{9}$$

$$\mathbb{E}[\bar{\ell}(0)]m = \mathbb{E}\left[\frac{1}{m}\sum_{e \in E} \ell_e(x)\right] m = \sum_{e \in E}\sum_{i=0}^{n} \Pr[i]\ell_e(i) \tag{10}$$

From (9) and (10), to prove $\mathbb{E}[\Phi(0)] = O(1) \times \mathbb{E}[\bar{\ell}(0)]m$ reduces to show that

$$\forall e \in E \text{ it holds } \sum_{i=0}^{n} \Pr[i] i \ell_e(i) - \sum_{i=0}^{n} \Pr[i]\ell_e(i) = O(1) \tag{11}$$

In turn, (11) equals:

$$\sum_{i=0}^{n} \Pr[i]\ell_e(i)(i-1) \leq \sum_{i=0}^{n} \Pr[i]\ell_e(i) i \leq \sum_{i=0}^{n} \Pr[i]\alpha^i i \approx e^{r(\alpha-1)}\alpha r = O(1) \tag{12}$$

□

However, `Greedy` may affect badly the initial distribution of bins, thus Lemma 3 may not hold for each $t > 0$.

It is helpful to consider the concurrent random process `Blind` (a simplified version of `Greedy`). At $t = 0$ throw randomly $n = O(m)$ users to m resources (`Blind`'s and `Greedy`'s initializations are identical). Initially, the load distribution has Binomial tail bounds from deviating from expectation $O(n/m) = O(1)$. During each round $t > 0$, `Blind` draws exactly 1 random user from each loaded resource (as Step 1 of `Greedy`). Let $n(t)$ be the subset of users drawn during round t. Round t ends by throwing at random these $|n(t)|$ users back into the m resources (then $|n(t)|$ users allowed by `Blind` to migrate is at least the $|\mathcal{A}(t)|$ ones allowed by `Greedy`, since no selfish criterion is required). Any resource is equally likely to receive any user, thus, `Blind` preserves per round $t > 0$ strong tail bounds from deviating from the constant expectation $O(n/m) = O(1)$ reminiscent to ones for $t = 0$. Therefore, Lemma 3 holds for each round $t > 0$ of `Blind`.

We show bellow that, as at round $t = 0$, similar strong tail bounds will hold for `Greedy` for each $t > 0$ with respect to the resources $\mathcal{L}_{\geq \alpha^{\nu_0}} = \{e \in E : \ell_e(f_e(t)) \geq \alpha^{\nu_0}\}$, with $\nu_0 = \lceil 2r \rceil + 1$, $r = \frac{n}{m} = O(1)$. Also let $\mathcal{L}_{<\alpha^{\nu_0}} = E \setminus \mathcal{L}_{\geq \alpha^{\nu_0}}$.

Lemma 4 *$\forall t \geq 0$ and any resource $e \in \mathcal{L}_{\geq \alpha^{\nu_0}}$ it holds $\mathbb{E}[f_e(t+1)] \leq f_e(t)$, with $\nu_0 = \lceil 2r \rceil + 1, r = \frac{n}{m} = O(1)$.*

Proof. Let $t \geq 0$ be any fixed round and let e be any fixed resource with $\ell_e(f_e(t)) \geq \alpha^{\nu}$. We observe:

$$\mathbb{E}[f_e(t+1)] = f_e(t) + \mathbb{E}[\text{\#users coming in } e \text{ in } t] - \mathbb{E}[\text{\#users leaving } e \text{ in } t] \qquad (13)$$

To establish the lemma, we show that if $\nu \geq \nu_0$ the expected number of users leaving e in t is no less than the expected number of users joining e in t.

Since $\ell_e(f_e(t)) \geq \alpha^{\nu}$, the α latency jump bound (Def. 1) implies that e can receive users only from resources in $E_{\geq \nu+1}(t) = \{j \in E : f_j(t) \geq \nu + 1\}$. In particular, resource e receives at most one player from each resource in $E_{\geq \nu+1}(t)$ with probability ϑ/m. Therefore,

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

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

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

By (13), it suffices to show that for every integer $\nu \geq \nu_0 = \lceil 2n/m \rceil + 1$, $|E_{\geq \nu+1}(t)| - |E_{\leq \nu-2}(t)| \leq 0$. Observe $|E_{\geq \nu+1}(t)| \leq |E_{\geq \nu-1}(t)|$ and $|E_{\leq \nu-2}(t)| = m - |E_{\geq \nu-1}(t)|$. Moreover,

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

$|E_{\geq \nu-1}(t)| \leq n/(\nu-1)$ 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 $\nu \geq \nu_0 = \lceil 2n/m \rceil + 1$. □

Lemma 5 *$\frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}} \mathbb{E}[\ell_e(f_e(t))] = \Theta(1)$ and $\frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}} \mathbb{E}[f_e(t)\ell_e(f_e(t))] = \Theta(1)$ with $\nu_0 = \lceil 2n/m \rceil + 1$.*

Proof. Consider an arbitrary resource e, at round $t_e > 0$ not in an (ε,α)-EQ, that leaves $\mathcal{L}_{<\alpha^{\nu_0}}$ and enters $\mathcal{L}_{\geq\alpha^{\nu_0}}$. By $\mathcal{L}_{<\alpha^{\nu_0}}$'s property, at round $t_e - 1$, e's latency was $< \alpha^{\nu_0}$. As e enters $\mathcal{L}_{\geq\alpha^{\nu_0}}$, Lemma 2 & Cor. 1 bound e's expected latency $\mathbb{E}[\ell_e(f_e(t_e))]$ as $\alpha^{\nu_0} \times (\alpha + \delta_\vartheta) = O(1)$, with $\delta_\vartheta = \frac{\varepsilon(\alpha-1)}{2}$ and α the latency jump bound (Def.1). At this point, the *Delta Method* [34] implies that for $\mathbb{E}[\ell_e(f_e(t_e))] = O(1)$ it holds

$$\mathbb{E}[\ell_e(f_e(t_e))] \approx \ell_e(\mathbb{E}[f_e(t_e)]) + \sigma^2 \times \frac{d^2\ell_e(x)}{dx^2}\,|_{x=\mathbb{E}[f_e(t_e)]} \tag{14}$$

with σ^2 the variance of $f_e(t_e)$ during round t_e and $\frac{d^2\ell_e(x)}{dx^2}\,|_{x=\mathbb{E}[f_e(t_e)]}$ the "curvature" of the latency ℓ_e, which (by Def. 1) is $\leq \alpha^{\mathbb{E}[f_e(t_e)]}\ln^2(\alpha)$. Using this, we show that $\forall t > t_e, \mathbb{E}[\ell_e(f_e(t))]$ will not become significantly higher than $\mathbb{E}[\ell_e(f_e(t_e))] < \alpha^{\nu_0} \times (\alpha + \delta_\vartheta) = O(1)$ attained at round t_e when e entered $\mathcal{L}_{\geq\alpha^{\nu_0}}$. To see this, observe in (14) that Lemma 4 (since $e \in \mathcal{L}_{\geq\alpha^{\nu_0}}$ and ℓ_e is increasing) gives: $\ell_e(\mathbb{E}[f_e(t_e+1)]) \leq \ell_e(f_e(t_e))$ and $\alpha^{\mathbb{E}[f_e(t_e+1)]}\ln^2(\alpha) < \alpha^{\mathbb{E}[f_e(t_e)]}\ln^2(\alpha)$. It only remains to show that during round $t_e + 1$ the variance σ^2 of variable $f_e(t_e+1)$ remains $O(1)$. To this end, observe that given $f_e(t_e)$ it holds: $f_e(t_e+1) = f_e(t_e) + (\sharp$users coming in $e) - (\sharp$users leaving $e) = f_e(t_e) + e_{in} - e_{out}$. As shown in the proof of Lemma 2, the newcomers e_{in} to e are statistically dominated by the binomial distribution $B(|\mathcal{A}(t_e+1)|, \vartheta/m)$, with $|\mathcal{A}(t_e+1)| \leq m$. Thus, $\sigma^2_{e_{in}} = O(1)$. Also, according to `Greedy` (Sec. 2.1), each resource looses 1 user with probability $\leq \vartheta/m$, and no user otherwise. Thus, for resource e we get $\sigma^2_{e_{out}} = O(1)$. We conclude that during round $t_e + 1$ for e holds: $\sigma^2 \leq \sigma^2_{e_{in}} + \sigma^2_{e_{out}} = O(1)$. We conclude that $\forall t > 0$ it holds:

$$\frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}} \mathbb{E}[\ell_e(f_e(t))] \;\leq\; \mathbb{E}[\bar{\ell}(t)] = \Theta(1) \tag{15}$$

Working similarly, we can show:

$$\frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}} \mathbb{E}[f_e(t)\ell_e(f_e(t))] \;\leq\; \mathbb{E}[\bar{\ell}(t)] \times \Theta(1) \tag{16}$$

□

Lemma 6 $\mathbb{E}[\Phi(t)] = \Theta(1) \times \mathbb{E}[\overline{\ell}(t)]m$

Proof.

$$\begin{aligned}
\mathbb{E}[\overline{\ell}(t)] &= \mathbb{E}\left[\frac{1}{m}\sum_{e\in E}\ell_e(x)\right] = \frac{1}{m}\sum_{e\in\mathcal{L}_{<\alpha^{\nu_0}}}\mathbb{E}[\ell_e(f_e(t))] + \frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}}\mathbb{E}[\ell_e(f_e(t))] \\
&= \frac{|\mathcal{L}_{<\alpha^{\nu_0}}|}{m}\alpha^{\nu_0} + \Theta(1) \qquad (17)
\end{aligned}$$

where the last equality is due to (15).

$$\begin{aligned}
\mathbb{E}[\Phi(t)] &= \mathbb{E}\left[\sum_{e\in\mathcal{L}_{<\alpha^{\nu_0}}}\sum_{i=1}^{f_e(t)}\ell_e(i) + \sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}}\sum_{i=1}^{f_e(t)}\ell_e(i)\right] \\
&\leq n\alpha^{\nu_0} + \sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}}\mathbb{E}[f_e(t)\ell_e(f_e(t))] = \left(r\alpha^{\nu_0} + \frac{1}{m}\sum_{e\in\mathcal{L}_{\geq\alpha^{\nu_0}}}\mathbb{E}[f_e(t)\ell_e(f_e(t))]\right)m \\
&= (r\alpha^{\nu_0} + \Theta(1))\, m \qquad (18)
\end{aligned}$$

where the last equality is due to (16).

The lemma follows by Expressions (17) and (18). □

2.6.3 Showing that $\sum_{i\in\mathcal{A}(t)}\mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] < -\frac{\varepsilon\vartheta\gamma^2}{4} \times \overline{\ell}(t)m$

Sketch of Case 1 and 2 below. A round is not (Def. 2) an (ε, α)-EQ if $\geq \varepsilon m$ resources have latency $\geq \alpha \times \overline{\ell}(t)$ or latency $\leq \frac{1}{\alpha} \times \overline{\ell}(t)$, with $\overline{\ell}(t)$ the average latency (Sec. 2.1) and $\alpha > 1$ the latency jump bound (Def. 1). These two options induce an expected potential drop studied in the corresponding Case 1 & 2 below. In each case, we write $\sum_{i\in\mathcal{A}(t)}\mathbb{E}[\Delta c_i(t)]$ as a portion of $-\overline{\ell}(t) \times m$. The underlying idea is: each migrant from an overloaded to an underloaded resource contributes to $\sum_{i\in\mathcal{A}(t)}\mathbb{E}[\Delta c_i(t)]$ her little portion of $\overline{\ell}(t)$ gain at hand. Roughly, we show that $\Omega(m)$ such migrations can boost the atomic gain $\overline{\ell}(t)$, when considered in the overall population $\mathcal{A}(t)$ of migrants, up to an $\Omega(1)$ portion of $\overline{\ell}(t) \times m$.

Case 1.

Let $\mathcal{U}_\gamma(t) = \{e \in E : \ell_e(f_e(t)) < (1-\gamma)\overline{\ell}(t)\}$ the *underloaded* and $\mathcal{O}(t) = \{e \in E : \ell_e(f_e(t)) \geq \alpha\overline{\ell}(t)\}$ the *overloaded* resources at round t, with parameter $\gamma \in (0, 1]$. Assume that $|\mathcal{O}(t)| \geq \varepsilon m$, that is, the state is not an (ε, α)-EQ (Def. 2), $\varepsilon \in (0, 1]$.

Fact 3 *If* $|\mathcal{O}(t)| \geq \varepsilon m$ *then*

1. $|\mathcal{U}_\gamma(t)| \geq \gamma m$,

2. *Each move from $\mathcal{O}(t)$ to $\mathcal{U}_\gamma(t)$ induces expected cost-decrease $\geq \frac{\gamma}{2}\bar{\ell}(t)$,*

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

with $\gamma = \frac{\varepsilon}{2}(\alpha-1)$, $\vartheta = \frac{\varepsilon}{4\alpha}$ the migration probability of `Greedy` *(Sec. 2.3), $\alpha > 1$ the latency jump bound (Def. 1), and $\bar{\ell}(t)$ the average latency (Sec. 2.1).*

Proof. For simplicity of notation, let $h = |\mathcal{O}(t)|$, $l = |\mathcal{U}_\gamma(t)|$, and $n = m - h - l$ the number of resources $\notin \mathcal{O}(t) \cup \mathcal{U}_\gamma(t)$ with latency $\in [(1-\gamma)\bar{\ell}(t), \alpha\bar{\ell}(t))$.

1. To reach a contradiction, assume $h \geq \varepsilon m$ and $l < \gamma m$, with $\gamma = \frac{\varepsilon}{2}(\alpha-1)$. Then $n > m - h - \gamma m$. The average latency (Sec. 2.1) becomes:

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

a contradiction. The 2nd ineq. holds because $\alpha\bar{\ell}(t)\,(h-\varepsilon m) \geq (1-\gamma)\bar{\ell}(t)\,(h-\varepsilon m)$, since $h \geq \varepsilon m$, $\alpha > 1$, and $\gamma > 0$. The 3rd ineq. holds because $\gamma = \frac{\varepsilon}{2}(\alpha-1)$ yields $\gamma(1-\varepsilon)+\gamma(1-\gamma) \leq \varepsilon(\alpha-1)$.

2. $\forall$ resource $e \in \mathcal{O}(t)$, a player migrates to $e' \in \mathcal{U}_\gamma(t)$ with probability $\geq \vartheta\gamma$ (see step (3) of `Greedy`-Sec. 2.3), where $\vartheta = \varepsilon/(4\alpha)$ and $\gamma = \frac{\varepsilon}{2}(\alpha-1)$ as shown in Fact 3(1.). In Lemma 2, set `Greedy`'s migration probability $\vartheta = \varepsilon/(4\alpha)$ obtaining that the corresponding δ_θ equals $\frac{\varepsilon}{4\alpha} = \frac{\delta_\theta}{\alpha(\alpha-1)} \Rightarrow \delta_\theta = \frac{\varepsilon(\alpha-1)}{4}$. The initial cost per user in $\mathcal{O}(t)$ is $\geq \alpha\bar{\ell}(t)$ and if she moves to $\mathcal{U}_\gamma(t)$, by Cor. 1, her expected cost will be $\leq (\alpha+\delta_\theta)\,(1-\gamma)\bar{\ell}(t) = \left(\alpha+\frac{\varepsilon(\alpha-1)}{4}\right)\left(1-\frac{\varepsilon(\alpha-1)}{2}\right)\bar{\ell}(t) = \left(\alpha+\frac{\varepsilon(\alpha-1)}{4}-\frac{\varepsilon\alpha(\alpha-1)}{2}-\frac{\varepsilon^2(\alpha-1)^2}{8}\right)\bar{\ell}(t) < \left(\alpha+\frac{\varepsilon(\alpha-1)}{4}-\frac{\varepsilon\alpha(\alpha-1)}{2}\right)\bar{\ell}(t) = \left(\alpha+\frac{\varepsilon(\alpha-1)}{2}\left(\frac{1}{2}-\alpha\right)\right)\bar{\ell}(t) = \left(\alpha+\gamma\left(\frac{1}{2}-\alpha\right)\right)\bar{\ell}(t)$ which is $< \left(\alpha-\frac{\gamma}{2}\right)\bar{\ell}(t)$ because $\gamma\left(\frac{1}{2}-\alpha\right) < -\frac{\gamma}{2}$ since $\alpha > 1$ & $\gamma \in (0,1)$, implying a cost decrease $\geq \frac{\gamma}{2}\bar{\ell}(t)$ per migrant, with $\gamma = \frac{\varepsilon}{2}(\alpha-1)$.

3. From Fact 3(2.), if k migrants switch from $\mathcal{O}(t)$ to $\mathcal{U}_\gamma(t)$ they induce an expected cost-drop $\geq \frac{\gamma}{2}\bar{\ell}(t)\times k$. Let $p_{\mathcal{O}\to\mathcal{U}_\gamma}[k|\bar{\ell}(t)]$ the probability that k such migrants appear, given $\bar{\ell}(t)$. The expectation $\sum_k kp_{\mathcal{O}\to\mathcal{U}_\gamma}[k|\bar{\ell}(t)]$ of such migrants is $\geq \varepsilon\vartheta\gamma m$, with $\vartheta = \frac{\varepsilon}{4\alpha}$, $\gamma = \frac{\varepsilon}{2}(\alpha-1)$ and α the latency jump bound (Def. 1). To see this, from each resource $e \in \mathcal{O}(t)$ with $|\mathcal{O}(t)| \geq \epsilon m$, exactly 1 player (step (1-2) of `Greedy`-Sec.2.3) may find an appealing resource $e' \in \mathcal{U}_\gamma(t)$, with probability $\geq \gamma$ since $|\mathcal{U}_\gamma(t)| \geq \gamma m$ due to Fact 3(1.). Then,

with probability ϑ she moves (step (3) of `Greedy`) to $e' \in \mathcal{U}_\gamma(t)$. Unconditioning on k, the expected cost-drop due to migrants switching from $\mathcal{O}(t)$ to $\mathcal{U}_\gamma(t)$ is

$$\geq \sum_{k\geq 0} \tfrac{\gamma}{2}\overline{\ell}(t) \times k p_{\mathcal{O}\to\mathcal{U}_\gamma}[k|\overline{\ell}(t)] \geq \tfrac{\gamma}{2}\overline{\ell}(t) \times \varepsilon\vartheta\gamma m = \tfrac{\varepsilon\vartheta\gamma^2}{2}\overline{\ell}(t)m \tag{19}$$

By (19) we finally prove (for Case 1) the result of this section:

$$\sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] \leq -\tfrac{\varepsilon\vartheta\gamma^2}{2} \times \overline{\ell}(t)m \tag{20}$$

□

Case 2.

Let $\mathcal{U}(t) = \{e \in E : \ell_e(f_e(t)) < \frac{1}{\alpha}\overline{\ell}(t)\}$ the *underloaded* and $\mathcal{O}_\gamma(t) = \{e \in E : \ell_e(f_e(t)) \geq (1+\gamma)\overline{\ell}(t)\}$ the *overloaded* resources at round t, with parameter $\gamma \in (0,1]$. Assume that $|\mathcal{U}(t)| \geq \varepsilon m$, that is, the state is not an (ε,α)-EQ (Def. 2), $\varepsilon \in (0,1]$.

Fact 4 *If $|\mathcal{U}(t)| \geq \varepsilon m$ then*

1. $\sum_{e\in\mathcal{O}_\gamma(t)} \ell_e(f_e(t)) > \gamma\overline{\ell}(t)m$,
2. *Each move from $\mathcal{O}_\gamma(t)$ to $\mathcal{U}(t)$ induces expected cost-decrease $\geq \frac{\gamma}{4}\ell_e(f_e(t))$,*
3. $\sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \leq -\frac{\vartheta\varepsilon\gamma^2}{4} \times m\overline{\ell}(t)$

with $\gamma = \frac{\varepsilon}{2\alpha}(\alpha-1)$, $\vartheta = \frac{\varepsilon}{4\alpha}$ the migration probability of `Greedy` *(Sec. 2.3), $\delta_\vartheta = \frac{\varepsilon(\alpha-1)}{2\alpha}$, $\alpha > 1$ the latency jump bound (Def. 1), and $\overline{\ell}(t)$ the average latency (Sec. 2.1).*

Proof. The line of thought is analogous to Fact 3.

1. Similarly to Fact 3(1.), we get: $\sum_{e\in O(t)} \ell_e(f_e(t)) > (1-\varepsilon/\alpha-(1-\varepsilon)(1+\gamma))\,\overline{\ell}(t)m = \gamma\,\overline{\ell}(t)m$, with $\gamma = \frac{\varepsilon(\alpha-1)}{2\alpha}$.
2. In Lemma 2 set $\vartheta = \frac{\varepsilon}{4\alpha}$ obtaining that the corresponding δ_ϑ 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 $e_1 \in \mathcal{O}_\gamma(t)$ is $\geq (1+\gamma)\,\overline{\ell}(t) = \left(1+\frac{\varepsilon(\alpha-1)}{2\alpha}\right)\overline{\ell}(t)$, with $\gamma = \frac{\varepsilon(\alpha-1)}{2\alpha}$. Each user in $e_1 \in \mathcal{O}_\gamma(t)$ if she moves to $e_2 \in \mathcal{U}(t)$, her expected latency becomes: $\leq \ell_{e_2}(f_{e_2}(t))(\alpha+\delta_\vartheta) = \ell_{e_2}(f_{e_2}(t))(\alpha+\alpha\frac{\gamma}{2}) = \ell_{e_2}(f_{e_2}(t))\alpha\left(1+\frac{\gamma}{2}\right) < \ell_{e_2}(f_{e_2}(t))\alpha\,(1+\gamma) < \frac{1}{\alpha}\overline{\ell}(t)\alpha\,(1+\gamma) = \overline{\ell}(t)\,(1+\gamma) \leq \ell_{e_1}(f_{e_1}(t))$ The 1st ineq. holds by Lemma 2. The 1st eq. stems by noticing $\delta_\theta = \frac{\alpha}{2}\gamma$. The 3rd. ineq. is due to $e_2 \in \mathcal{U}(t)$. The last ineq. follows from $e_1 \in \mathcal{O}_\gamma(t)$. Therefore $\ell_{e_1}(f_{e_1}(t)) \geq$

$\alpha(1+\gamma)\ell_{e_2}(f_{e_2}(t)) \Rightarrow \ell_{e_2}(t) \le \frac{1}{\alpha(1+\gamma)}\ell_{e_1}(t)$. Then, if user migrates to e_2 its expected latency will be $\le \frac{\alpha+\delta_\theta}{\alpha(1+\gamma)}\ell_{e_1}(t) = \frac{\alpha+\frac{\alpha}{2}\gamma}{\alpha(1+\gamma)}\ell_{e_1}(t) = \frac{1+\frac{\gamma}{2}}{1+\gamma}\ell_{e_1}(t)$. Its expected cost decrease will be $\left(1-\frac{1+\frac{\gamma}{2}}{1+\gamma}\right)\ell_{e_1}(t) = \frac{\gamma}{2(1+\gamma)}\ell_{e_1}(t) \ge \frac{\gamma}{4}\ell_{e_1}(t)$, since $\gamma = \frac{\varepsilon(\alpha-1)}{2\alpha} < 1$.

3. Similarly to Fact 3(3.), since $|\mathcal{U}(t)| \ge \varepsilon m$, any player on $e \in \mathcal{O}_\gamma(t)$ moves to $e' \in \mathcal{U}(t)$ with probability $\ge \vartheta\varepsilon$. Recall from Fact 4(2.), that such a move induces expected cost decrease $\ge \frac{\gamma}{4}\ell_e(f_e(t))$, with $\gamma = \frac{\varepsilon(\alpha-1)}{2\alpha}$. Combining this with Fact 4(1.), we obtain that the overall expected cost-drop due to migrants leaving resources $\mathcal{O}_\gamma(t)$ and entering $\mathcal{U}(t)$ in round t is at least:

$$\vartheta\varepsilon \times \frac{\gamma}{4}\sum_{e\in\mathcal{O}(t)} \ell_e(f_e(t)) > \vartheta\varepsilon \times \frac{\gamma}{4} \times \gamma\,\overline{\ell}(t)m > \frac{\vartheta\varepsilon\gamma^2}{4}\,\overline{\ell}(t)m \tag{21}$$

By (21) we finally prove (for Case 2) the result of this section:

$$\sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] \le -\tfrac{\vartheta\varepsilon\gamma^2}{4} \times \overline{\ell}(t)m \tag{22}$$

□

2.6.4 Proof of Theorem 2.

Here we combine the results in Section 2.6.1, 2.6.2 and 2.6.3 and prove Theorem 2. From Section 2.6.1 we get $\mathbb{E}[\Delta\Phi(t)] \le \sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)] \le 0$. Fix an arbitrary average latency (Sec. 2.1) $\overline{\ell}(t)$ where `Greedy` is not on an (ε,α)-EQ. Facts 3&4 in Sec. 2.6.3 yield:

$$\mathbb{E}[\Delta\Phi(t)|\overline{\ell}(t)] \le \sum_{i\in\mathcal{A}(t)} \mathbb{E}[\Delta c_i(t)|\overline{\ell}(t)] < -\tfrac{\varepsilon\vartheta\gamma^2}{4} \times \overline{\ell}(t)m \tag{23}$$

Consider the space of all realizations $\overline{\ell}(t)$ not in an (ε,α)-EQ. Let $\Pr[\overline{\ell}(t)]$ the probability to obtain such a realization $\overline{\ell}(t)$. Removing the conditional on $\overline{\ell}(t)$, Expression (23) becomes:

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

From Lemma 6 the above becomes: $\mathbb{E}[\Delta\Phi(t)] \le -\frac{\vartheta\varepsilon\delta^2}{4} \times \Theta(1) \times \mathbb{E}[\Phi(t)]$.

3 Approximate Equilibria in Congestion Games with Coalitions

In this section, we investigate the rate of convergence to approximate pure NE in an "organised" setting, where the players are pre-partitioned into static coalitions, and in each round, only the coalition with the largest improvement in its total delay is allowed to move. Hence, we consider a setting where selfish moves are coordinated and sequential. Nevertheless, each coalitional move contains a certain amount of concurrency, since it may change the strategies of many players simultaneously.

For single-commodity linear congestion games with static coalitions, we establish an upper bound on the rate of convergence to approximate NE that is polynomial in the number of coalitions k and the ratio r of the largest coalition's size to the smallest coalition's size, and logarithmic in the potential of the initial state. In the special case that the number of coalitions is constant and the coalitions are almost equisized, we obtain a logarithmic upper bound on the convergence time to approximate NE.

3.1 Model and Preliminaries

A *congestion game with coalitions* consists of a set of identical players $N = [n]$[4] partitioned into k coalitions $\{C_1, \dots, C_k\}$, a set of resources $E = \{e_1, \dots, e_m\}$, a strategy space $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$ for each player $i \in N$, and a non-negative and non-decreasing latency function $\ell_e : \mathbb{N} \mapsto \mathbb{R}_+$ 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 $\Sigma^{n_1} \times \cdots \times \Sigma^{n_k}$. 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 of the largest coalition's size to the smallest coalition's size.

A pure strategy $s_j \in \Sigma^{n_j}$ determines a (pure) strategy $s_j^i \in \Sigma$ for every player $i \in C_j$. 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 due to C_j in s_j is $f_e(s_j) = |\{i \in C_j : e \in s_j^i\}|$. 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 $\sigma \in \Sigma$ in state s is $\ell_\sigma(s) = \sum_{e \in \sigma} \ell_e(f_e(s))$.

[4]For every integer $n \geq 1$, we let $[n] \equiv \{1, \dots, n\}$.

The individual 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 [15, Theorem 2] and then computing a min-cost flow.

A state s is a *pure 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[5]. For some $\varepsilon \in (0, 1)$, 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 *$\varepsilon$-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 in s.

If the current state is not an ε-Nash equilibrium, there may be many coalitions with ε-moves available. In the following, we consider the (sequential) *largest improvement ε-Nash dynamics*, where the coalition that moves next is the one whose best response move is an ε-move and results in the largest improvement in its total delay. Formally, for a state s that is not an ε-Nash equilibrium, the coalition that moves next in the largest improvement ε-Nash dynamics, is a coalition C_j such that (i) for any other coalition C_i,

$$\max_{s'_j \in \Sigma^{n_j}} \{\tau_j(s) - \tau_j(s_{-j}, s'_j)\} \geq \max_{s'_i \in \Sigma^{n_i}} \{\tau_i(s) - \tau_i(s_{-i}, s'_i)\},$$

and (ii) $\max_{s'_j \in \Sigma^{n_j}} \{\tau_j(s) - \tau_j(s_{-j}, s'_j)\} > \varepsilon\tau_j(s)$ (ties are resolved arbitrarily).

3.2 Convergence to Approximate Equilibria

To bound the convergence time of the largest improvement ε-Nash dynamics, we use the following potential function:

$$\Phi(s) = \tfrac{1}{2} \sum_{e \in E} [f_e(s)\ell_e(f_e(s)) + \sum_{j=1}^{k} f_e(s_j)\ell_e(f_e(s_j))] \tag{24}$$

[19, 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 static coalitions, the (sequential) largest improvement ε-Nash dynamics converges to an ε-Nash equilibrium in a number

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

of steps that is polynomial in k and r and logarithmic in the potential of the initial state. Hence for network congestion games, where a coalition's best response can be computed in polynomial time by a min-cost flow computation, an ε-Nash equilibrium can be computed in polynomial time. Furthermore, in the special case that the number of coalitions is constant and the coalitions are almost equisized, i.e. when $k = \Theta(1)$ and $r = \Theta(1)$, the largest improvement ε-Nash dynamics converges in a logarithmic number of steps.

Theorem 5 *In a single-commodity linear congestion game with n players divided into k coalitions, the sequential 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 of the largest coalition's size to the smallest coalition's size.*

Proof. The outline of the proof is similar to that of [8, Theorem 3.4], which holds for symmetric congestion games only. However, coalitions may be of different size, in which case the game is not symmetric. Hence, we have to extend the technique of [8] 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 [8].

Let $\{C_1, \ldots, C_k\}$ be a set of coalitions, and let $s = (s_j)_{j\in[k]}$ be a state 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 the best response of C_i 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 in s, 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 [19, Theorem 6]). 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 coalition of maximum total delay 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,

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

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

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

References

[1] H. Ackermann, H. Roeglin, and B. Voecking. On the impact of combinatorial structure on congestion games. In *FOCS*, 2006.

[2] B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *STOC*, 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] O. Babaoglu, G. Canright, A. Deutsch, G. A. Di Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemani, and T. Urnes. Design Patterns from Biology for Distributed Computing. *ACM Transactions on Autonomous and Adaptive Systems*, 1(1):26–66, 2006.

[5] P. Berenbrink, T. Friedetzky, L. A.Goldberg, P. Goldberg, Z. Hu, and R. Martin. Distributed selfish load balancing. In *SODA*, 2006.

[6] 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.

[7] 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.

[8] S. Chien and A. Sinclair. Convergece to Approximate Nash Equilibria in Congestion Games. In *SODA*, 2007.

[9] G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *STOC*, 2005.

[10] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *STACS*, 2006.

[11] G. Cybenko. Load balancing for distributed memory multiprocessors. *J. Parallel Distrib. Comput.*, 7:279–301, 1989.

[12] M. Dahlin. Interpreting Stale Load Information. *IEEE Transactions on Parallel and Distributed Systems*, 11(10), Oct 2001.

[13] E. Even-Dar and Y. Mansour. Fast convergence of selfish rerouting. In *SODA*, 2005.

[14] Eyal Even-Dar, Alex Kesselman, and Yishay Mansour. Convergence Time to Nash Equilibria. In *ICALP*, 2003.

[15] A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *STOC*, 2004.

[16] S. Fischer, H. Räcke, and B. Vöcking. Fast convergence to wardrop equilibria by adaptive sampling methods. In *STOC*, 2006.

[17] S. Fischer and B. 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.

[18] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *TCS*, 348:226–239, 2005.

[19] D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic Congestion Games among Coalitions. In *ICALP*, 2006.

[20] Bhaskar Ghosh and S. Muthukrishnan. Dynamic load balancing in parallel and distributed networks by random matchings. In *Proc. of the 6th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA).*, pages 220–225, 1994.

[21] M. X. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS'05*.

[22] P. W. Goldberg. Bounds for the convergence rate of randomized local search in a multi-player load-balancing game. In *PODC*, 2004.

[23] T. Grenager, R. Powers, and Y. Shoham. Dispersion games: general definitions and some specific learning results. In *AAAI*, 2002.

[24] A. Hayrapetyan, É. Tardos, and T. Wexler. The Effect of Collusion in Congestion Games. In *STOC*, 2006.

[25] S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *AAAI*, 2005.

[26] 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.

[27] 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.

[28] J. F. Kurose and K.W. Ross. *Computer Networking, a top down approch featuring the Internet*. Addison-Wesley Longman (3rd ed.), 2004.

[29] Lavy Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.

[30] V. Mirrokni and A. Vetta. Convergence issues in competitive games. In *APPROX'04*.

[31] Michael Mitzenmacher. How useful is old information? *IEEE Trans. Parallel Distrib. Syst.*, 11(1):6–20, 2000.

[32] D. Monderer and L. Shapley. Potential Games. *Games& Econ. Behavior*, 14:124–143, 1996.

[33] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[34] G. W. Oehlert. A Note on the Delta Method. *The American Statistician*, 46(1):27–29, 1992.

[35] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. on Net.*, 1(5):510–521, 1993.

[36] J. Rexford. *Handbook of optimization in Telecommunications, chapter Route optimization in IP networks*. Kluwer Academic Publishers, 2005.

[37] R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[38] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. To be published by MIT Press. http://www.ssc.wisc.edu/ whs/book/index.html.

[39] W. H. Sandholm. Potential Games with Continuous Player Sets. *Journal of Economic Theory*, 97:81–108, 2001.

[40] J. W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.