Convergence to Approximate Nash Equilibria in Congestion Games†

Steve Chien‡ Alistair Sinclair§

Abstract

We study the ability of decentralized, local dynamics in non-cooperative games to rapidly reach an approximate Nash equilibrium. For symmetric congestion games in which the edge delays satisfy a "bounded jump" condition, we show that convergence to an ε-Nash equilibrium occurs within a number of steps that is polynomial in the number of players and ε^{-1}. This appears to be the first such result for a class of games that includes examples for which finding an exact Nash equilibrium is PLS-complete, and in which shortest paths to an exact equilibrium are exponentially long. We show moreover that rapid convergence holds even under only the apparently minimal assumption that no player is excluded from moving for arbitrarily many steps. We also prove that, in a generalized setting where players have different "tolerances" ε_i that specify their thresholds in the approximate Nash equilibrium, the number of moves made by a player before equilibrium is reached depends only on his associated ε_i, and not on those of the other players. Finally, we show that polynomial time convergence still holds even when a bounded number of edges are allowed to have arbitrary delay functions.

1 Introduction

The emerging field of algorithmic game theory has led to a fundamental re-examination, from a computational perspective, of the classical concept of Nash equilibrium [20]. Much of this activity has focused on understanding the structure of Nash equilibria (as expressed, notably, in the "price of anarchy," see e.g. [22, 25, 24]) and the computational complexity of finding them (see, e.g., [10, 7, 4]). Considerably less is understood about the question of whether selfish players, acting in a decentralized fashion, actually arrive at a Nash equilibrium in a reasonable amount of time. This would seem to be a central consideration in the computational study of Nash equilibria.

In this paper we address this question in the general arena of *congestion games*. A congestion game is an n-player game in which each player's strategy consists of a set of resources, and the cost of the strategy depends only on the number of players using each resource, i.e., the cost takes the

†A preliminary version of this paper appeared in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 2007.

‡Microsoft Research Silicon Valley Campus, 1065 La Avenida, Mountain View CA 94043, U.S.A. Email: schien@microsoft.com.

§Computer Science Division, University of California, Berkeley CA 94720-1776, U.S.A. Email: sinclair@cs.berkeley.edu. Supported in part by NSF ITR grant CCR-0121555. Work done in part while this author was visiting Microsoft Research Silicon Valley Campus.

form $\sum_e d_e(f(e))$, where $f(e)$ is the number of players using resource e, and d_e is a non-negative increasing function. A standard example is a *network* congestion game on a directed graph, in which each player must select a path from some source to some destination, and each edge has an associated "delay" function that increases with the number of players using the edge. In what follows, we shall use the terminology of edges and delays even though we will always be discussing general (non-network) congestion games.

Congestion games have attracted a good deal of attention, partly because they capture a large class of routing and resource allocation scenarios, and not least because they are known to possess *pure* Nash equilibria [23]. Thus unlike general games, whose Nash equilibria may involve mixed (i.e., randomized) strategies for the players, congestion games always have a Nash equilibrium in which each player sticks to a single strategy. Further, in congestion games, the natural decentralized mechanism known as the "Nash dynamics", in which at each step some player switches her strategy to a better alternative, is guaranteed to converge to a pure Nash equilibrium. The question then is the following: *Starting from an arbitrary initial state, does the Nash dynamics converge rapidly?*

The work of [10] provides a devastating negative answer, even for symmetric[†] congestion games: the problem of finding a Nash equilibrium is *PLS-complete* [14], and therefore as difficult as that of finding a local optimum in any local search problem with efficiently computable neighborhoods. Moreover, there are examples of games and initial strategies such that the shortest path to an equilibrium in the Nash dynamics is exponentially long in the number of players n. Thus if we want a notion of Nash equilibrium that is selfishly and efficiently realizable, the best we can hope for is some kind of approximation. (Indeed, given the recent spate of hardness results for finding exact equilibria in most classes of games by *any* algorithmic means [10, 7, 4], it seems inevitable that attention will now shift to approximation.)

Accordingly, we say that a state s (i.e., a collection of strategies for the players) is an *ε-Nash equilibrium* if no player can improve her cost by more than a factor of ε by unilaterally changing her strategy. This definition has intuitive appeal, for example, if one imagines charging players a percentage of their current cost for the privilege of changing strategy.[‡] Given this definition, we introduce a natural modification of the Nash dynamics called the *ε-Nash dynamics*, which permits only *ε-moves*, i.e., moves that improve the cost of the player by a factor of more than ε. Clearly ε-Nash equilibria correspond to fixed points of this dynamics. Our goal is to investigate under what circumstances the ε-Nash dynamics does in fact converge rapidly to an ε-Nash equilibrium.

To make the ε-Nash dynamics concrete, we assume that among multiple players with ε-moves available, at each step a move is made by the player with the largest incentive to move; i.e., the player who can make the largest relative improvement in cost (with ties broken arbitrarily). This is a minimal coordination mechanism that seems natural in our context; however, as we shall see later, our results hold even with *no* coordination under only a basic liveness assumption.

In order to state our results we need one further notion. For any $\alpha \geq 1$, we say that an edge in a congestion game satisfies the *α-bounded jump* condition if its delay function satisfies $d_e(t+1) \leq \alpha d_e(t)$ for all $t \geq 1$. We will think of α as being a constant, or at most polynomially bounded in n. The bounded jump condition means that when a new player is added to an edge,

[†]A *symmetric* game is one in which the allowed strategies of all the players are the same.

[‡]An alternative notion of approximate equilibrium (see, e.g., [8, 10, 15, 16]) is based on an *additive* error of ε, rather than the relative error we use here. We would argue that our definition is equally natural, and indeed more in line with approximation guarantees in Computer Science and also with the notion of price of anarchy in game theory [22].

the cost to all players using that edge increases by at most a factor of α. This condition is rather weak (see below); in particular, an edge with $d_e(t) = \alpha^t$ satisfies the α-bounded jump condition.

We are now ready to state our first main result, which says that in any symmetric congestion game with bounded jumps, the ε-Nash dynamics converges rapidly to an ε-Nash equilibrium. This is apparently the first such result for such a broad class of (atomic) congestion games, and in particular for a class that contains PLS-complete examples.

Theorem 1.1 *In any symmetric congestion game with n players in which all edges satisfy the α-bounded jump condition, the ε-Nash dynamics converges from any initial state in $\lceil n\alpha\varepsilon^{-1}\log(nC)\rceil$ steps, where C is an upper bound on the cost of any player.*

The proof of this theorem relies on two fundamental principles. First, the existence of an "exact" potential function [23], whose decrease under any move reflects exactly the improvement in cost of the moving player. And second, the fact that, under the bounded jump condition, any player can emulate the move of any other with at most an α-factor overhead. This ensures that every move of the dynamics decreases the potential function by an $\frac{\varepsilon}{\alpha n}$ factor.

We now briefly discuss the bounded jump condition. Firstly, as we show later (Section 3.1), the hardness results mentioned above for finding exact equilibria carry over to symmetric games in which all edges have α-bounded jumps. Secondly, we claim that the bounded jump condition is a reasonable assumption in practice, and is similar to conditions imposed in other quantitative studies of transient behavior (e.g., the "bounded relative slope" of [11] or "bounded slope" of [3]); it is also much weaker than the polynomial bounds typically used in studies of the price of anarchy [1, 6]. Thirdly, it is questionable how much sense it makes to talk about "symmetric" congestion games without such a condition. This is because of the trick in [10] (see Section 3.1 below) for making any congestion game symmetric by adjoining to the strategies of each player p_i a special edge e_i whose delay is small for one player and huge for more than one player. This effectively divides the strategies into sets, one player per set, and is equivalent to the original game up to a relabeling of the players. Thus, if we could prove Theorem 1.1 without the bounded jump condition, we would get rapid convergence for *all* congestion games. The bounded jump condition can be seen as expressing an alternative, stronger notion of symmetry: no player can effectively "lock out" another by using a resource whose cost would explode if an additional player were to use it.

Next, we investigate the role of the order of player moves in ensuring rapid convergence. Recall that our ε-Nash dynamics assumed that moves are made by players with the largest available relative cost improvement. We show that our rapid convergence result, Theorem 1.1, is robust under any reasonable variation of the ε-Nash dynamics, including the "largest gain" dynamics (the player who moves is one who can gain the largest *absolute* cost improvement by an ε-move) and the "heaviest first" dynamics (the player who moves is one with largest current cost among those with an ε-move available). Our most convincing illustration is the "unrestricted" dynamics, in which an adversary may specify which player is allowed to move at each step, subject only to the basic liveness condition that no player is prevented from moving for arbitrarily many steps. This includes, for example, the "round-robin" scheme where in each round all players are selected to move according to some fixed permutation.

Theorem 1.2 *In any symmetric n-player congestion game whose edges satisfy the α-bounded jump condition, any ε-Nash dynamics in which all players are given an opportunity to move within each time interval of length T converges from any initial state in $\lceil\frac{n(\alpha+1)}{\varepsilon(1-\varepsilon)}\log(nC)\rceil T$ steps.*

We then go on to consider a natural generalization of the ε-dynamics to "heterogeneous" players, each of whom has an individual tolerance value $\varepsilon = \varepsilon_i$. Thus player p_i has an incentive to move only if she can improve her cost by a factor of more than ε_i. A straightforward generalization of Theorem 1.1 bounds the number of steps of this dynamics in terms of the smallest tolerance value $\varepsilon_{\min} = \min_i \varepsilon_i$. However, it is natural to ask if one can say more; in particular, if some player has a relatively large value of ε_i (and thus is very "tolerant"), can this player be forced to move very many times because of other, less tolerant players in the system? We prove an intriguing result along these lines. We show that the number of time steps at which a player with tolerance ε_i will be "unhappy" (i.e., will have an ε_i-move available) is essentially $O(n\alpha\varepsilon_i^{-1}\log(nC))$, irrespective of the ε_j-values of the other players! Thus highly intolerant players are not able to force others to move frequently.

Theorem 1.3 *Let $\varepsilon_{\max} < 1$ be the maximum value of ε_i among all players p_i. Then for any value $\varepsilon > 0$, there are at most $\lceil \frac{n\alpha}{\varepsilon(1-\varepsilon_{\max})}\log(nC)\rceil$ times at which some player p_j with $\varepsilon_j \geq \varepsilon$ will be able to move before the ε-Nash dynamics converges.*

Finally, we investigate the extent to which the bounded jump assumption can be relaxed. Specifically we prove the following:

Theorem 1.4 *In the setting of Theorem 1.1, the ε-Nash dynamics converges in* $\mathrm{poly}(n, \alpha, \varepsilon^{-1}, \log nC)$ *steps even if a fixed number of edges violate the α-bounded jump condition.*

Thus rapid convergence of the dynamics is still assured even if a constant number of edges have arbitrarily large jumps in their delay functions. In light of our earlier discussion of symmetry, we can view this as a step towards extending our results to *asymmetric* games: Theorem 1.4 allows us to have any constant number of "classes" of players, with each class selecting its strategies from a specific set. (By contrast, a symmetric game has just one class, while a general asymmetric game may have as many as n classes.) The proof of Theorem 1.4 is rather more technical, and involves the introduction of what we call "reduced games" involving certain subsets of the players.

The remainder of the paper is organized as follows. In Section 1.1 we give a brief summary of related work. In Section 2 we set notation and define our central concepts. Section 3 proves our basic convergence result, Theorem 1.1, including versions based on different orders of player moves. In Section 4 we prove Theorem 1.2, establishing rapid convergence for the unrestricted dynamics. Section 5 considers heterogeneous players and proves Theorem 1.3. Finally, in Section 6 we extend the analysis to the case of a constant number of edges with arbitrary delay functions and prove Theorem 1.4. We conclude with some open problems in Section 7.

1.1 Related work

Fabrikant, Papadimitriou and Talwar [10] systematically studied the complexity of finding Nash equilibria in congestion games; in particular they showed that finding a Nash equilibrium in symmetric congestion games is PLS-complete (and thus hard for local search). They also gave a polynomial time (global) algorithm for the case of symmetric *network* congestion games, but this algorithm says nothing about convergence of local dynamics.

Convergence questions similar to those in the present paper (i.e., the Nash dynamics, or some simple local learning algorithm for the players) have been investigated by other authors in various contexts. There are a number of results on "load-balancing" games, which are restricted congestion games in which each strategy consists of just a single edge (or "machine"), but which may be generalized to allow either player-specific cost functions [17] or weights on the players [8, 9, 13]. Milchtaich [17], Even-Dar et al. [8] and Goldberg [13] establish polynomial time convergence for versions of the Nash dynamics to (exact or approximate) Nash equilibria in these games, while Even-Dar and Mansour [9] consider a more complex dynamics in which all players move concurrently according to a certain rerouting mechanism. Kearns and Mansour [15] give polynomial time global and local algorithms that find additive ε-approximate equilibria for "large-population" games under a "bounded influence" assumption; however, this assumption appears not to hold for the general multiple-resource congestion games we consider here.

Recent papers by Fischer, Räcke and Vöcking [11] and Blum, Even-Dar and Ligett [3] consider congestion games at a similar level of generality to ours, each with some version of a "bounded (relative) slope" assumption that is analogous to bounded jumps. However, these papers analyze the *non-atomic* setting where the number of players is taken to infinity (the so-called "Wardrop traffic model"). Despite its apparent similarity, the non-atomic case is actually quite different from our discrete setting; for example, in the non-atomic case Nash equilibria can be computed in polynomial time [2, 10]. Fischer *et al.* [11] establish polynomial bounds on the rate of convergence to approximate Nash equilibria (under a different notion of approximation) of a concurrent dynamics with moves based on "adaptive sampling", while Blum *et al.* [3] give polynomial bounds when players use no-regret online learning algorithms.

We mention also two recent developments for more general games. Goemans, Mirrokni and Vetta [12, 18] study convergence of Nash dynamics not to an (approximate) Nash equilibrium but instead to a "sink equilibrium", for wider classes of games for which pure equilibria need not exist. They quantify the rate of convergence in various cases, and also the quality of the resulting solution as measured by a global utility function, rather than player-specific costs. And a subexponential time non-local algorithm for finding an approximate Nash equilibrium in general games with a fixed number of players and explicitly presented strategies was given by Lipton, Markakis and Mehta [16].

Finally we note that Theorem 1.1 is reminiscent of, and partially inspired by, recent work of Orlin, Punnen, and Schulz [21], who show how to find an ε-approximate local minimum for any problem in PLS. However, our setting differs from theirs in two crucial respects. Firstly, their algorithm finds an ε-local minimum of the potential function, which is not necessarily an ε-Nash equilibrium. Secondly, their algorithm would require ostensibly selfish players to somehow be aware of the value of the global potential function and to limit their actions based on this knowledge.

2 Background

Congestion games. A *game* consists of a finite set of players $\{p_1, \ldots, p_n\}$, each of which is assigned a finite set of *strategies* S_i and a cost function $c_i : S_1 \times \cdots \times S_n \to \mathbb{N}$ that he wishes to minimize. A game is called *symmetric* if all of the S_i are identical. A *state* $s = (s_1, \ldots, s_n) \in S_1 \times \cdots \times S_n$ is any combination of strategies for the players. A state s is a *pure Nash equilibrium* if for all players p_i, $c_i(s_1, \ldots, s_i, \ldots, s_n) \leq c_i(s_1, \ldots, s'_i, \ldots, s_n)$ for all $s'_i \in S_i$; thus at a Nash equilibrium, no player can improve his cost by unilaterally changing his strategy. It is well known that, while every (finite) game has a *mixed* Nash equilibrium[§], not every game has a pure Nash equilibrium.

[§]In a mixed Nash equilbrium, a player's strategy can be any probability distribution over available strategies, and no individual player can improve his expected cost by choosing another probability distribution.

We will focus on the class of games known as *congestion games*, where players' costs are based on the shared usage of a common set of resources, which we shall call *edges* $E = \{e_1, \ldots, e_m\}$. A player's strategy set $S_i \subseteq 2^E$ is an arbitrary collection of subsets of E; his strategy $s_i \in S_i$ will therefore be a subset of E. Each edge $e \in E$ has an associated nondecreasing delay function $d_e : \{1, \ldots, n\} \to \mathbb{N}$; if t players are using the edge e, they will each incur a cost of $d_e(t)$. As a result, in a state $s = (s_1, \ldots, s_n)$, the cost of player p_i is $c_i(s) = \sum_{e \in s_i} d_e(f_s(e))$, where $f_s(e)$ is the number of players using edge e under s (i.e., $f_s(e) = |\{j : e \in s_j\}|$).

Existence of potential functions and pure Nash equilibria. Congestion games possess several appealing characteristics, including the existence of an *exact* potential function. This function is defined as

$$\phi(s) = \sum_{e \in E} \sum_{t=1}^{f_s(e)} d_e(t), \tag{1}$$

and has the property that if player p_i shifts strategy from s_i to s_i', the change in ϕ exactly mirrors the change in the player's cost: i.e., $\phi(s) - \phi(s') = c_i(s) - c_i(s')$ [23].

An important consequence of this is the observation that, if we follow an iterative process where at each step one player changes strategy to lower his cost (a *Nash dynamics*), then the potential function ϕ will decrease until it reaches a local minimum, which must be a pure Nash equilibrium. However, this does not provide a bound on the number of such player moves required to reach a pure Nash equilibrium in a congestion game. Indeed, as mentioned in the Introduction, it has been shown [10] that there exist (symmetric) congestion games in which the number of player moves required to go from one state to any pure Nash equilibrium is exponentially large.

Approximate Nash equilibria and ε-Nash dynamics. We define an *ε-Nash equilibrium* as a state in which no player has more than an ε-incentive to move:

Definition 2.1 *For $\varepsilon \in [0, 1)$, a state $s = (s_1, \ldots, s_n) \in S_1 \times \cdots \times S_n$ is an* ε-Nash equilibrium *if for all players p_i, $c_i(s_1, \ldots, s_i', \ldots, s_n) \geq (1 - \varepsilon)c_i(s_1, \ldots, s_i, \ldots, s_n)$ for all $s_i' \in S_i$.*

We complement this definition with that of the *ε-Nash dynamics*, where we require that a player may make only *ε-moves*, or moves that improve his cost by a factor of more than ε; i.e., if player p_i moves from s_i to s_i' then $c_i(s_1, \ldots, s_i', \ldots, s_n) < (1-\varepsilon)c_i(s_1, \ldots, s_i, \ldots, s_n)$. Clearly when no further ε-moves are possible, the players have reached an ε-Nash equilibrium. Further, for concreteness, we stipulate that if more than one player has an ε-move available, a player whose relative gain is largest will be the one that moves. Thus a move is made by a player p_i who maximizes $\frac{c_i(s) - c_i(s_1, \ldots, s_i', \ldots, s_n)}{c_i(s)}$. This choice seems the most natural, and unless otherwise stated we shall assume it throughout the paper. However, as we shall demonstrate later our results are not sensitive to this choice, and hold for a wide variety of other natural variations of the dynamics; in particular, they hold for the *unrestricted* ε-Nash dynamics, which allows an adversarial order of player moves so long as every player is offered the chance to move every so often.

Bounded jumps. We say that an edge e satisfies the *α-bounded jump condition* if its delay function satisfies $d_e(t + 1) \leq \alpha d_e(t)$ for all $t \geq 1$, for some value $\alpha \geq 1$. In our applications, we shall think of α as being constant or at most polynomially bounded in n. This still allows delay functions as large as α^t, and is therefore a much weaker restriction than the Lipschitz condition (see, e.g., [10]), which requires that the delay functions be linearly bounded. (Note that one consequence of our definition is that $d_e(1) > 0$; otherwise $d_e(t) = 0$ for all t, and e is essentially irrelevant.) As we will see later (Section 3.1), even for symmetric congestion games with the bounded jump condition on all edges, finding a Nash equilibrium can be PLS-complete; thus in this sense, bounded jumps are not a major restriction on the power of congestion games.

3 The basic convergence theorem

The main purpose of this section is to show the following, which is a restatement of Theorem 1.1 from the Introduction:

Theorem 3.1 *In any symmetric congestion game with n players in which all edges satisfy the α-bounded jump condition, the ε-Nash dynamics converges from any initial state in $\lceil n\alpha\varepsilon^{-1}\log(nC)\rceil$ steps, where C is an upper bound on the cost of any player.*

(Note that here and elsewhere, our bound is undefined for the case of exact Nash equilibria, i.e., when $\varepsilon = 0$.)

Before giving the proof, we sketch the basic structure of the argument, which will be used repeatedly in the paper. The key observation is that after polynomially many moves by players with high costs, we must necessarily reach an ε-Nash equilibrium. For this purpose we use the exact potential function ϕ defined in equation (1). Suppose player p_i, with current cost $c_i(s) \geq \frac{\phi(s)}{\beta}$, makes an ε-move; this move must reduce c_i, and hence ϕ, by more than $\frac{\varepsilon\phi(s)}{\beta}$. After at most about $\beta\varepsilon^{-1}\log\phi_{\max}$ such steps, where $\phi_{\max}$ is the initial value of the potential function, we must have reached an ε-Nash equilibrium. Since $\phi(s) \leq \sum_i c_i(s)$, if the highest-cost player moves then we may take $\beta = n$ and we are done. The main challenge, then, is to show that high-cost players move reasonably frequently, and are not blocked by low-cost players whose moves do not significantly decrease ϕ.

In light of the above discussion, the following lemma will be the main tool in the proof.

Lemma 3.2 *In a symmetric congestion game in which every edge has α-bounded jumps, if in the ε-Nash dynamics with state s the next move is made by player p_i, then $c_j(s) \leq \alpha c_i(s)$ for all j.*

Proof: Suppose player p_i moves from s_i to s'_i, taking the game from state $s = (s_1, \dots, s_n)$ to $s' = (s_1, \dots, s'_i, \dots, s_n)$. Consider an arbitrary player p_j, and the resulting state if p_j, rather than p_i, had adopted s'_i; denote this state $s'' = (s_1, \dots, s_i, \dots, s''_j = s'_i, \dots, s_n)$.

Since p_i moved and not p_j, we can conclude that p_j's relative gain for this move is at most p_i's relative gain, regardless of whether this is an ε-move for p_j. (If it is an ε-move, then p_i's relative gain must be at least as large by the definition of the dynamics; if it is not an ε-move, then p_j's relative gain is at most ε while p_i's relative gain is more than ε.) Thus we have

$$\frac{c_j(s) - c_j(s'')}{c_j(s)} \leq \frac{c_i(s) - c_i(s')}{c_i(s)}. \tag{2}$$

Now let us compare the cost p_i pays for adopting s'_i, namely $c_i(s')$, with how much p_j would have paid for the same strategy, namely $c_j(s'')$. For each edge $e \in s'_i$, either p_i is already occupying it before the move ($e \in s_i$), or not. In the former case, p_j may have to pay as much as $d_e(f_s(e)+1)$ to use e, while p_i only pays $d_e(f_s(e))$; by the bounded jump assumption, these differ by at most a factor of α. In the latter case, p_i pays $d_e(f_s(e)+1)$ and p_j pays at most the same amount. Summing over all edges $e \in s'_i$, we obtain $c_j(s'') \leq \alpha c_i(s')$.

Combining this with inequality (2), we obtain $\frac{c_j(s)-\alpha c_i(s')}{c_j(s)} \leq \frac{c_i(s)-c_i(s')}{c_i(s)}$, from which we can see that $c_j(s) \leq \alpha c_i(s)$, as required. $\square$

Proof of Theorem 3.1: Lemma 3.2 guarantees that every time any player (say p_i) moves, the cost of that player is at least $\frac{1}{\alpha}$ times the largest cost of any player. Since for any state s we have that $\phi(s) \leq \sum_j c_j(s)$, then $c_i(s) \geq \frac{1}{\alpha n}\phi(s)$. But under a move of p_i taking the game from state s to s', the decrease in the potential function is $\phi(s) - \phi(s') = c_i(s) - c_i(s') > \varepsilon c_i(s) \geq \frac{\varepsilon}{\alpha n}\phi(s)$. Thus at each move ϕ must decrease by a factor of more than $\frac{\varepsilon}{\alpha n}$. But since ϕ is non-negative integer-valued, there can be at most $\lceil n\alpha\varepsilon^{-1} \log \phi_{\max} \rceil$ such decreases, where $\phi_{\max}$ is the initial value of the potential function. Since clearly $\phi_{\max} \leq nC$, we are done. $\square$

Remark: Note that we may replace $\log(nC)$ in Theorem 3.1 by $\log(\phi_{\max}/\phi_{\min})$, where $\phi_{\max}$, $\phi_{\min}$ are upper and lower bounds respectively on the possible values of the potential function.

3.1 PLS-completeness of bounded jump games

We complement Theorem 3.1 by observing that the class of congestion games with bounded jumps on all edges includes examples for which it is PLS-complete to find an exact Nash equilibrium; indeed, for such games the shortest path to an exact equilibrium in the Nash dynamics can be exponentially long, while Theorem 3.1 shows that an ε-equilibrium is reached in a polynomial number of steps.

Proposition 3.3 *The problem of finding a Nash equilibrium in symmetric congestion games satisfying the α-bounded jump condition with $\alpha = 2$ is PLS-complete.*

Proof: We follow the chain of reductions in Theorem 3 of [10], but with some modifications to the delay functions. The starting point is POSNAE3FLIP: given an instance of not-all-equal-3SAT with weights on the clauses and only positive literals, find a truth assignment such that the total weight of all satisfied clauses cannot be improved by flipping the value of a single variable. This problem is known to be PLS-complete [26]. In [10, Theorem 3(i)] this is reduced to the problem of finding a Nash equilibrium in a congestion game as follows. For each 3-clause c there are two edges, e_c and e'_c, with delay functions $d(1) = d(2) = 0$ and $d(t) = w_c$ for $t > 2$. There is one player for each variable x, and the player has two strategies: one contains all the e_c for clauses c that contain x, and the other contains all the e'_c for the same clauses. Any Nash equilibrium of this game corresponds to a local optimum of POSNAE3FLIP. Now observe that, since both strategies of any player contain the same number of edges, with the same delay functions, the Nash equilibria will not be affected if we add a constant to the delay functions of all edges; i.e., the delay becomes $d(1) = d(2) = w_c$ and $d(t) = 2w_c$ for $t > 2$.

Finally, following [10, Theorem 3(ii)] we can reduce this game to a symmetric game by adjoining to both the strategies of each player x a new edge e_x with delay function $d_{e_x}(1) = D$ and $d_{e_x}(t) = 2D$ for $t > 1$, where $D = 2mw_{\max} + 1$. (Here m is the number of clauses, and $w_{\max}$ the maximum weight of a clause.) Clearly, in any Nash equilibrium of the symmetric game with the same number of players and all strategies available to all, exactly one strategy from each pair in the original game must in fact be selected. Thus we have arrived at a symmetric congestion game in which the delay function of every edge satisfies the α-bounded jump condition for $\alpha = 2$, and for which finding a Nash equilibrium is PLS-complete. Moreover, as observed in [10], the reductions inherit from [26] the property that the length of a shortest path to an equilibrium may be exponentially long. This completes the proof of the Proposition. □

Remark: In the above construction, the value $\phi_{\max}/\phi_{\min}$ can also be seen to be at most 3, so by the remark following the proof of Theorem 3.1 the convergence time of the ε-Nash dynamics is $O(n\varepsilon^{-1})$.

We inject one caveat into the above discussion: altering the delay functions on the edges as we did has no effect on the Nash equilibria, but may have a significant effect on the ε-equilibria. Thus the approximate equilibria found in Theorem 3.1 for the modified game may bear little relationship to those for the original game. Our only purpose here was to demonstrate that Theorem 3.1 can apply to games whose exact Nash equilibria are hard to locate.

3.2 Variations on the dynamics

We now discuss some variations on the ε-Nash dynamics, and show that Theorem 3.1 still holds in these cases. Thus the rapid convergence to an ε-equilibrium guaranteed by the theorem does not depend crucially on allowing the player with largest relative gain to move. In the next section, we will show how to dispense with coordination altogether.

Largest gain dynamics. Define the *largest gain* ε-Nash dynamics as that in which, at each step, among all players with an ε-move available, the one that moves is one whose (absolute) cost improvement is greatest. We show that Theorem 3.1 still holds under this dynamics:

Theorem 3.4 *Theorem 3.1 continues to hold under the largest gain ε-Nash dynamics.*

Proof: Consider any move in the dynamics that takes the game from state s to state s'. It suffices to show that this causes the potential function ϕ to drop by a factor of at least $\frac{\varepsilon}{\alpha n}$; i.e., $\phi(s) - \phi(s') \geq \frac{\varepsilon}{\alpha n}\phi$. To see this, let p_i be the player that moves, and consider any other player p_j. We examine two cases: either p_j has an ε-move available, or p_j does not.

In the first case, p_i's absolute improvement must be at least $\varepsilon c_j(s)$, since p_j could have improved by at least $\varepsilon c_j(s)$, but p_i was given priority by virtue of having a larger absolute gain. Thus $\phi(s) - \phi(s') \geq \varepsilon c_j(s)$ for p_j with ε-moves available.

In the second case, let s'' be the resulting state if p_j were to move to s'_i instead of p_i. By the same argument as in the proof of Lemma 3.2, we have that $c_j(s'') \leq \alpha c_i(s')$. Since $c_i(s') < (1-\varepsilon)c_i(s)$, but $(1-\varepsilon)c_j(s) \leq c_j(s'')$, we can conclude that $c_j(s) < \alpha c_i(s)$. Hence $\phi(s) - \phi(s') > \varepsilon c_i(s) > \frac{\varepsilon}{\alpha} c_j(s)$ for all p_j with no ε-moves available.

Combining the two cases, we have that $\phi(s) - \phi(s') \geq \frac{\varepsilon}{\alpha} c_j(s)$ for all players p_j. Since at least one player p_j has cost $c_j(s) \geq \frac{1}{n}\phi(s)$, we obtain $\phi(s) - \phi(s') \geq \frac{\varepsilon}{\alpha n}\phi(s)$, as required. □

Heaviest first dynamics. If at each step, among all players with an ε-move available, we allow a player with largest current cost to move, we arrive at the *heaviest first* ε-Nash dynamics. We can show that this version of the dynamics also leads to rapid convergence:

Theorem 3.5 *Theorem 3.1 continues to hold under the heaviest first ε-Nash dynamics.*

Proof: It suffices to show that Lemma 3.2 still holds under this dynamics. At any given step, let p_i be the player that moves, $s = (s_1, \ldots, s_n)$ be the current state, and $s' = (s_1, \ldots, s'_i, \ldots, s_n)$ be the state after p_i moves. We wish to show that $c_j(s) \leq \alpha c_i(s)$ for all players p_j.

Fix j and define s'' to be the state that results if p_j were to move to s'_i instead of p_i. Since p_j did not do so, we conclude that either $c_j(s) \leq c_i(s)$, in which case p_i had priority over p_j in the heaviest first dynamics and the lemma is satisfied (with equality when $\alpha = 1$), or $c_j(s) > c_i(s)$ but also $(1-\varepsilon)c_j(s) \leq c_j(s'')$.

To handle this latter case, note first that, as in the proof of Lemma 3.2, we have $c_j(s'') \leq \alpha c_i(s')$. Moreover, as above, since p_i makes a move from s_i to s'_i, we must have $c_i(s') < (1-\varepsilon)c_i(s)$. Putting all of this together, we get

$$(1-\varepsilon)c_j(s) \leq c_j(s'') \leq \alpha c_i(s') < (1-\varepsilon)\alpha c_i(s),$$

so we can conclude that $c_j(s) < \alpha c_i(s)$ as required. □

4 The unrestricted dynamics

So far our Nash dynamics, while essentially decentralized, has assumed some minimal coordination mechanism whereby players with the largest incentive move first. In this section, we show that our results still hold (with a modest penalty in convergence time) when this coordination is removed. We consider what is in a sense the most liberal possible dynamics, which we call the *unrestricted* ε-Nash dynamics. In this dynamics, players may move in an arbitrary order (which may even be under the control of an adversary), subject only to a minimal "liveness" condition. This condition just says that every player must be *given an opportunity to move*[¶] within a bounded amount of time; without such a condition, one or more players could be "locked out" for arbitrarily long and we could not expect to bound the rate of convergence. More formally, the unrestricted dynamics is specified by a sequence $q_1, q_2, \ldots$, where each q_t denotes a player; at step t, player q_t is given the opportunity to move, and actually makes a move if he has an ε-move available. (Otherwise nothing happens at step t.) The sequence (q_t) may be adaptive (i.e., q_t may depend in an arbitrary way on the past, the current state etc.) We require only that, for some constant T, in each interval of the sequence of length T every player p_i appears at least once. A natural dynamics satisfying this condition is the "round-robin" dynamics, where in each round all players are selected to move according to some fixed permutation.

We now show that Theorem 3.1 holds for the unrestricted dynamics, with a slightly looser bound on the time until convergence. Note that this seemingly stronger result does not in fact imply polynomial convergence in the setting of Theorem 3.1 (or any of its variants above) because the original ε-Nash dynamics may not satisfy the liveness condition. (Some player may never be in a position of having a move with the largest relative gain.)

Theorem 4.1 *In any symmetric n-player congestion game whose edges satisfy the α-bounded jump condition, any ε-Nash dynamics in which every player is given an opportunity to move within each time interval of length T converges from any initial state in $\lceil \frac{n(\alpha+1)}{\varepsilon(1-\varepsilon)} \log(nC) \rceil T$ steps, where C is an upper bound on the cost of any player.*

Before proving the theorem, we state and prove an important lemma that allows us to relate improvements in the potential function to a change in the cost of a player, even when the player does not move for many steps. As can be seen from the proof, the lemma is not specific to the unrestricted dynamics but holds for any variant of the ε-Nash dynamics. Moreover, the proof makes

[¶]Note that this player need not have an ε-move available! Thus the adversary may attempt to always select players who cannot move.

no use of the bounded jump property; in fact, we shall apply the lemma to games with unbounded jumps in Section 6.

Lemma 4.2 *Let $c_i(s)$ be the cost of player p_i in state s, and let $c_i(s')$ be p_i's cost in a future state s' in which p_i has not moved. Then $\phi(s) - \phi(s') \geq \varepsilon(c_i(s) - c_i(s'))$.*

Proof: Note that $c_i(s) - c_i(s') = \sum_{e \in s_i} d_e(f_s(e)) - d_e(f_{s'}(e))$, and that the only positive contributions to this sum are from those edges e in which $f_s(e) > f_{s'}(e)$, i.e., edges that other players have vacated. For each such edge e, the first player p_j to give up e must have had cost at least $d_e(f_s(e))$ at the time, and hence improved the potential function by at least $\varepsilon d_e(f_s(e))$. Therefore the total improvement to the potential function can be bounded as follows:

$$\phi(s) - \phi(s') \geq \sum_{e: f_s(e)) > f_{s'}(e)} \varepsilon d_e(f_s(e)) \geq \varepsilon(c_i(s) - c_i(s')). \quad \square$$

Proof of Theorem 4.1: It is sufficient to show that during any interval in which every player is given an opportunity to move, the potential function ϕ must decrease by at least $\frac{\varepsilon(1-\varepsilon)}{(\alpha+1)n}\phi^0$, where ϕ^0 is the value of the potential function at the beginning of the interval. Denote the states during this interval as $s^0, s^1, \ldots, s^T$; note that successive states in this sequence need not be distinct, as the player licensed to move may not in fact be able to move.

Let p_h be the player with largest cost in s^0; let $t \geq 0$ be the first time during this interval in which p_h is given the chance to move. We analyze two cases:

Case (i): At time t, p_h has an ε-move available. From Lemma 4.2, we are guaranteed that $\phi(s^0) - \phi(s^t) \geq \varepsilon(c_h(s^0) - c_h(s^t))$, and thus after p_h moves ϕ will have improved by at least $\varepsilon c_h(s^0) \geq \frac{\varepsilon}{n}\phi(s^0)$. So in this case the claim above holds.

Case (ii): At time t, p_h does not have an ε-move available. In this case, we observe that at time t, we must have

$$c_h(s^t) \leq \frac{\alpha}{1-\varepsilon} c_i(s^t) \quad \text{for all other players } p_i. \tag{3}$$

If not, for any player p_i violating this condition, p_h can make an ε-move by simply adopting p_i's strategy at an overall cost of at most $\alpha c_i(s^t)$. Now note that at least one player must actually move in the interval $[0, \ldots, T]$; otherwise, we are already at an ε-Nash equilibrium. Suppose on the one hand that the first player to move, say p_i, does so at time $t' > t$, i.e., after p_h has been given a chance to move. Then we have $c_h(s^t) = c_h(s^0)$ and $c_i(s^{t'}) = c_i(s^t) = c_i(s^0)$, and combining this with (3) we obtain $c_i(s^{t'}) \geq \frac{1-\varepsilon}{\alpha} c_h(s^0)$. The improvement to the potential function caused by this move is then at least $\varepsilon \frac{1-\varepsilon}{\alpha} c_h(s^0) \geq \frac{\varepsilon(1-\varepsilon)}{\alpha} \frac{\phi(s^0)}{n}$.

Now suppose on the other hand that some player moves before time t, and let p_i be the last such player to move, this move taking place at time $t' < t$. We claim that

$$c_h(s^t) \leq \frac{\alpha}{1-\varepsilon} c_i(s^{t'}). \tag{4}$$

To see this, note from (3) that p_i must satisfy this condition at time t, and hence also immediately after the move (which is the last before time t); and since the move can only decrease his cost he must satisfy the condition at time t' also. Now the change in ϕ up to time t is bounded as follows:

$$\begin{aligned} \phi(s^0) - \phi(s^t) &\geq \max\left\{\varepsilon c_i(s^{t'}),\, \varepsilon\left(c_h(s^0) - c_h(s^t)\right)\right\} \\ &\geq \varepsilon \max\left\{\frac{1-\varepsilon}{\alpha} c_h(s^t),\, c_h(s^0) - c_h(s^t)\right\}. \end{aligned}$$

In the first line, the first item in the maximum is the improvement gained by p_i for his move, while the second item follows from Lemma 4.2. The second line comes from inequality (4) above.

Finally, this last expression is minimized when $c_h(s^t) = \frac{\alpha}{\alpha+1-\varepsilon} c_h(s^0)$, and thus the potential function must decrease by at least $\frac{\varepsilon(1-\varepsilon)}{\alpha+1-\varepsilon} c_h(s^0) \geq \frac{\varepsilon(1-\varepsilon)}{\alpha+1} \frac{\phi(s^0)}{n}$. This concludes the analysis of case (ii) and hence the proof. $\square$

5 Heterogeneous players

We now generalize our previous setting by allowing each player p_i to have her own value $\varepsilon = \varepsilon_i$ that specifies her "tolerance" for unhappiness. Thus whereas one relaxed player may be content with $\varepsilon = 0.1$, another may be more particular and demand $\varepsilon = 0.001$. Accordingly, given individual player tolerances $\varepsilon_i \in [0, 1)$, we extend our definition of approximate Nash equilibrium as follows:

Definition 5.1 *For* $\boldsymbol{\varepsilon} = (\varepsilon_i) \in [0, 1)^n$*, a state* $s = (s_1, \ldots, s_n) \in S_1 \times \cdots \times S_n$ *is an* $\boldsymbol{\varepsilon}$-Nash equilibrium *if for all players* p_i*,* $c_i(s_1, \ldots, s_i', \ldots, s_n) \geq (1 - \varepsilon_i) c_i(s_1, \ldots, s_i, \ldots, s_n)$ *for all* $s_i' \in S_i$*.*

The $\boldsymbol{\varepsilon}$-Nash dynamics is also extended in the obvious way. For definiteness, we will go back to our original ε-Nash dynamics of Section 3 in which the player with largest relative gain moves.

By modifying the proof of Theorem 3.1, it can be shown that this dynamics converges in $O(n\alpha\varepsilon_{\min}^{-1}(1 - \varepsilon_{\max})^{-1} \log(nC))$ steps, where $\varepsilon_{\min} = \min_i \varepsilon_i$ and $\varepsilon_{\max} = \max_i \varepsilon_i$. However, it is natural to ask if one can say more; in particular, if some player has a relatively large value of ε_i (and thus is very "tolerant"), can this player be forced to move very many times because of other, less tolerant players in the system?

We now prove an intriguing result along these lines. We show that the number of time steps at which a player with tolerance ε_i will be "unhappy" (i.e., will have an ε_i-move available) is essentially $O(n\alpha\varepsilon_{\min}^{-1} \log(nC))$, irrespective of the ε_j-values of the other players. Thus highly intolerant players are not able to force others to move frequently.

Theorem 5.2 *Let* $\varepsilon_{\max} < 1$ *be the maximum value of* ε_i *among all players* p_i*. Then for any value* $\varepsilon > 0$*, there are at most* $\lceil \frac{n\alpha}{\varepsilon(1-\varepsilon_{\max})} \log(nC) \rceil$ *times at which some player* p_j *with* $\varepsilon_j \geq \varepsilon$ *will be able to move before the* $\boldsymbol{\varepsilon}$*-Nash dynamics converges.*

Proof: Consider a state $s = (s_1, \ldots, s_n)$ in which a player p_j with $\varepsilon_j \geq \varepsilon$ has an ε_j-move available. It suffices to show that the decrease in potential function ϕ is at least $\frac{\varepsilon_j(1-\varepsilon_{\max})}{\alpha n} \phi(s)$.

Let p_i be the player who actually moves from state s, and let $s' = (s_1, \ldots, s_i', \ldots, s_n)$ be the resulting new state. The largest relative gain dynamics implies that $\phi(s') - \phi(s) > \varepsilon_j c_i(s)$.

Now let p_h be the player with largest cost in s. If $p_h = p_i$ then we are done immediately, since $c_h(s) \geq \frac{\phi(s)}{n}$. Otherwise, let s'' be the state that results from s if p_h moves instead of p_i and takes p_i's new strategy s_i'. As in the proof of Theorem 3.1, we have $c_h(s'') \leq \alpha c_i(s') < \alpha(1 - \varepsilon_j) c_i(s)$. Since p_h does not actually move from s, either (1) p_h moving to s'' is not an ε_h-move for p_h; or (2) the relative gain that p_h gets from such a move is no more than the relative gain p_i gets from its move. We analyze these two cases separately.

In the first case, we have that $c_h(s) - c_h(s'') \leq \varepsilon_h c_h(s)$. Since $c_h(s'') < \alpha(1 - \varepsilon_j) c_i(s)$, we can conclude that $c_h(s) - \alpha c_i(s) < \varepsilon_h c_h(s)$, and therefore $c_i(s) \geq \frac{1-\varepsilon_h}{\alpha} c_h(s)$. The change in potential function is then at least $\phi(s) - \phi(s') > \frac{\varepsilon_j(1-\varepsilon_h)}{\alpha n} \phi(s)$.

In the second case, we have $\frac{c_h(s)-c_h(s'')}{c_h(s)} \leq \frac{c_i(s)-c_i(s')}{c_i(s)}$, or $\frac{c_h(s'')}{c_h(s)} \geq \frac{c_i(s')}{c_i(s)}$. Again, since $c_h(s'') < \alpha c_i(s')$, we conclude that $c_h(s) \leq \alpha c_i(s)$. Thus $\phi(s) - \phi(s') \geq \frac{\varepsilon_j}{\alpha n}\phi(s)$. □

Remarks:

1. The above theorem includes an additional factor $(1 - \varepsilon_{\max})^{-1}$, and thus says little when some ε_i is very close to 1. We believe that this is a technical artifact of the proof. Moreover, if ε_i is very close to 1 then the corresponding player, p_i, will move only when he is able to reduce his cost to essentially zero in one move; clearly this is not a scenario of great practical interest.

2. Unlike our other results, Theorem 5.2 is somewhat sensitive to the choice of ε-Nash dynamics. For example, while it also holds for the largest gain dynamics, it can fail for the heaviest first dynamics.

5.1 Wait times with heterogeneous players

Theorem 5.2 leads to the following natural question: Given that all players with ε-values larger than any specific ε will collectively be able to make an ε-move a limited number of times, can we place a polynomial upper bound on the last time at which any such player will be able to move?

We show that the answer is no. In particular, we construct a symmetric bounded-jump congestion game with $n+1$ players, as well as a starting state, in which n of the players, $p_1, \ldots, p_n$, each have $\varepsilon_i = 0$ and have to solve a PLS-complete problem before the final player, p_*, who has some positive $\varepsilon_* > 0$, is able to move.

We do this by embedding the PLS-complete problem Posnae3flip into our game in a manner similar to that outlined in the proof of Proposition 3.3. Given an instance of Posnae3flip on n variables, the embedded game has a player p_i for each variable x_i, with tolerance $\varepsilon_i = 0$. Again, there are two edges e_c and e'_c for each clause c; both these edges have delay function $d(1) = d(2) = w_c$, and $d(3) = 2w_c$. Each player p_i has the same two strategies as before, one containing all edges e_c for each clause c containing x_i, and one containing all edges e'_c for the same clauses. Assume that this instance of Posnae3flip is such that there exist states from which any path to a Nash equilibrium is exponentially long; we start the game with the p_i in one of these states. We will denote this initial state by s^0, and subsequent states $s^1, s^2, \ldots$ until the game reaches a Nash equilibrium s^N. Note that each edge in the embedded game has at most three players on it at each step; let σ^r denote the set of clauses c for which either e_c or e'_c has exactly three players in state s^r.

We now introduce our new player p_* with $\varepsilon_* > 0$, who has two strategies available. One strategy s_* is p_*'s initial strategy, and contains only a single new edge e_*, while the other strategy s'_* contains all previous edges e_c and e'_c for all clauses c.

We now set the delay function for e_* and extend the delay functions for the e_c and e'_c. For each clause $c \in \sigma^N$, we set both $d_{e_c}(4)$ and $d_{e'_c}(4)$ equal to $d_{e_c}(3) = 2w_c$, while for all $c \notin \sigma^N$, we set both $d_{e_c}(4)$ and $d_{e'_c}(4)$ equal to $d_{e_c}(3) + \gamma$, where γ is a small constant (say, 2). Finally, we set $d_{e_*}(1) = \beta$, where with some foresight we set β to be $\left\lceil \frac{\sum_c 3w_c}{1-\varepsilon_*} + 1 \right\rceil$.

Since the original n players p_i, if left to themselves, will take an exponential number of steps to reach their own Nash equilibrium, we will be done if we can show that the new player p_* cannot make an ε_*-move until this Nash equilibrium is reached, but will make such a move as soon as this happens. In other words, if we denote by $c_*(s^r \cup s'_*)$ the cost that p_* would pay for using strategy s'_* when the other players are in state s^r, then we need to verify that $c_*(s^r \cup s'_*) \geq (1 - \varepsilon_*)\beta$ for all $r < N$, while $c_*(s^N \cup s'_*) < (1 - \varepsilon_*)\beta$.

To see this, observe that at any time step $0 \le r \le N$, the cost that p_* would pay for his other strategy s'_* is $c_*(s^r \cup s'_*) = \sum_{c \in \sigma^r}(d_{e_c}(4) + d_{e_c}(1)) + \sum_{c \notin \sigma^r}(d_{e_c}(3) + d_{e_c}(2)) = \sum_c 3w_c + \sum_{c \in \sigma^r \setminus \sigma^N} \gamma$. Since $\sigma^r \not\subseteq \sigma^N$ for all $r < N$ (if $\sigma^r \subseteq \sigma^N$, then $\phi(s^r) \le \phi(s^N)$ where ϕ is the potential function for the original game, before the addition of player p_*), we have that $c_*(s^r \cup s'_*) \ge \sum_c 3w_c + \gamma \ge (1-\varepsilon_*)\beta$ for all $r < N$, and $c_*(s^N \cup s'_*) = \sum_c 3w_c < (1-\varepsilon_*)\beta$.

Finally, we apply the same trick as in the proof of Proposition 3.3 to make the game symmetric. We observe that, as in that proof, the resulting game satisfies the α-bounded jump condition with $\alpha = 2$.

6 Congestion games with unbounded jumps

We now investigate what happens when we relax the requirement that every edge in the congestion game satisfies the bounded jump condition. Our goal is to prove polynomial convergence of the ε-Nash dynamics even when we allow any constant number of edges to have arbitrary delay functions. Specifically, we will prove the following theorem, which is a more precise reformulation of Theorem 1.4 in the Introduction.

Theorem 6.1 *In any symmetric congestion game with n players in which all but k edges satisfy the α-bounded jump condition, the ε-Nash dynamics converges from any initial state in at most $\lceil n\alpha\varepsilon^{-1}\log(nC)\rceil^{2^k}$ steps, where C is an upper bound on the cost of any player.*

Before presenting the proof we sketch some of the main ideas. First, consider the simple case in which there is only one edge, e^*, that is not α-bounded. Our previous analysis fails since Lemma 3.2 no longer holds: now, when player p_i makes a move, it is no longer true that another player p_j can match that move with an α-factor penalty. To overcome this obstacle, we introduce the concept of a *reduced game*; the reduced game in state s consists of exactly those players p_i whose strategies s_i contain e^*. When a player in the reduced game gives up e^*, or a player outside the reduced game makes a move, then we know that ϕ must drop by the usual $\frac{\varepsilon}{\alpha n}$, since any player could emulate this move; call these moves *good* moves. To observe that this must happen frequently, note that the reduced game is itself a game with α-bounded jumps! Hence by our previous results the reduced game can only continue for a small number of steps before reaching an ε-equilibrium.

To extend the analysis to an arbitrary number k of exceptional edges, we apply the above idea recursively. Starting from the global game on all n players, we build a nested sequence of reduced games by repeatedly excluding from the next game the heaviest remaining player p_h, along with all players whose exceptional edges are a subset of those held by p_h. Thus each successive reduced game contains fewer players than its predecessor, and is indexed by a subset of exceptional edges that is not a subset of any previous reduced game; clearly this sequence of games has length at most 2^k. With this structure in place, we can show that each move in the global game causes a good move *in one of the reduced games*, and hence only a limited number of these moves can happen before the reduced games reach equilibrium; at this point a good global move must occur. There are several technical details to handle; in particular, when a good move is made in one of the reduced games, all subgames after it in the sequence have to be redefined; and the analysis requires an extension of Lemma 4.2 that applies in this context.

6.1 Proof of Theorem 6.1

We now give the details of the proof. Recall that our context is a congestion game in which there are k "exceptional" edges that do not satisfy the α-bounded condition, each of which may have a different, arbitrary nondecreasing delay function. We will denote the set of exceptional edges by $E^* = \{e_1^*, \ldots, e_k^*\} \subseteq E$.

We begin with the following extension of Lemma 4.2.

Lemma 6.2 *Fix an arbitrary state s and player p_i, and let $Q \subseteq E^*$ be the subset of exceptional edges held by p_i in s. Let s' be the state immediately after one of the following first occurs: (1) a player whose exceptional edges are a subset of Q makes a move; or (2) a player makes a move that results in his exceptional edges being a subset of Q. Then $\phi(s) - \phi(s') > \frac{\varepsilon c_i(s)}{\alpha}$.*

This lemma says the following. Consider the first time a player (say, p_j) makes a move that p_i could also make without adopting any additional exceptional edges that p_j currently holds. Then over the intervening time interval, the decrease in the potential function has at least the same guarantee as if p_i himself had moved, up to a factor of α.

Proof: Let p_j denote the player that makes the move resulting in the game reaching s', and let s'' be the state just before p_j makes this move; note that p_i has not moved by this point. Then by Lemma 4.2 (which as observed earlier holds even for games with unbounded jumps) we have

$$\phi(s) - \phi(s'') \geq \varepsilon(c_i(s) - c_i(s'')). \tag{5}$$

Now consider the move of player p_j that takes the state from s'' to s' (i.e., p_j's strategy changes from s_j'' to s_j'); this decreases p_j's cost by more than a factor of ε, so $\phi(s'') - \phi(s') > \varepsilon c_j(s'')$. To get a lower bound on this decrease, we consider the cost that p_i would pay had he, rather than p_j, made the move to s_j'; to compute this, we need to consider how much it would cost p_i to obtain the edges in $s_j' \setminus s_i$. From the statement of the lemma, we have two cases to consider.

In the first case, the set of exceptional edges held by p_j in s_j'' (before the move) is a subset of those held by p_i. Thus for each exceptional edge $e \in s_j'$, p_i would only have to pay at most the same cost as p_j. Meanwhile, for each bounded jump edge $e \in s_j'$, p_i may have to pay as much as a factor of α more than p_j. Summing over both classes of edges yields that p_i would pay at most $\alpha c_j(s')$ to adopt s_j' from the state s''. However, since p_i did not do so, we know that his relative gain can be at most that of p_j, so that

$$\frac{c_i(s'') - \alpha c_j(s')}{c_i(s'')} \leq \frac{c_j(s'') - c_j(s')}{c_j(s'')}.$$

This implies $c_i(s'') \leq \alpha c_j(s'')$, and hence $\phi(s'') - \phi(s') > \frac{\varepsilon c_i(s'')}{\alpha}$. Combining this with (5), we conclude that $\phi(s) - \phi(s') > \frac{\varepsilon c_i(s)}{\alpha}$.

In the second case, the set of exceptional edges held by p_j in s_j' (after the move) is a subset of those held by p_i. Again, p_i pays at most the same cost as p_j for these edges, and the rest of the argument proceeds as in the first case. □

We now formally define the concept of a *nested sequence of reduced games*, as indicated in the proof sketch. Given a congestion game with k exceptional edges E^* in a particular state s, we can describe it as a nested sequence of smaller congestion games as follows:

- Choose a player p_i; let Q denote the set of exceptional edges held by p_i. Identify all players p_j (including p_i) whose sets of exceptional edges are a subset of Q; place these players in the top level of the sequence.

- Recursively perform this same operation on the remaining players until no players are left.

It is easy to see that when this procedure is completed, the players will have been partitioned into $\ell \leq 2^k$ levels, with each level being associated with a particular subset of E^*. If we number the levels from 0 to $\ell - 1$, with the top level being 0, we see that the subsets Q_i associated with each level i form a partial order; namely if $Q_i \subseteq Q_j$, then $i \leq j$. A player will belong to the first level i at which his set of exceptional edges Q is a subset of Q_i. W.l.o.g. we assume that $Q_{\ell-1} = E^*$.

Further, for any level i, we can define the *reduced game* at level i, denoted G_i, as the game consisting of all players at levels $i, \ldots, \ell-1$, equipped with delay functions and a potential function ϕ_i that includes only these players. More formally, if P_i is the set of players in G_i, then the delay function on each edge is $d_e^{(i)}(t) = d_e(\hat{f}_s^{(i)}(e) + t)$, where $\hat{f}_s^{(i)}(e)$ is the number of players not in G_i that utilize e. The set of strategies for each player is as in the original game, excluding strategies whose exceptional edges are a subset of Q_j for some $j < i$. The potential function ϕ_i is defined analogously, with $\phi_i(s) = \sum_{e \in E} \sum_{t=1}^{f_s(e) - \hat{f}_s^{(i)}(e)} d_e^{(i)}(t)$.

We make use of this nested sequence as follows. Given a congestion game and an initial state s, we create a sequence of games by choosing, at each level, the *heaviest* player among all remaining players (with ties broken arbitrarily). This defines a particular initial sequence of games. After each move in the global game, we reorganize this sequence as follows:

- Consider the set of exceptional edges that the moving player held before his move and after his move, and determine which of these sets falls into a higher level (lower value of i) in the sequence. Let G_i be the reduced game corresponding to this level.

- Consider all players in the reduced game G_i, and recreate the sequence from this point down, again choosing the heaviest player to determine each level. We will say that the games G_j, $j \geq i$ have been *(re)initialized* at this point.

An important point to realize here is that, while the set of players at *level* i changes after reinitialization, the set of players in the *game* G_i remains the same, and so the potential function ϕ_i for G_i retains its meaning. All games (and potential functions) for $j > i$ are created again from scratch.

Lemma 6.3 *When a game G_i is the highest game to be reinitialized after a move, its corresponding potential function ϕ_i will have decreased by more than $\varepsilon \frac{\gamma_i}{\alpha}$ since its last initialization, where γ_i was the largest cost of any player in game G_i at the time of the last initialization of G_i.*

Proof: From the description above, it is clear that all moves that have occurred since the last initialization of G_i have been made by players in G_i; a move by a player in G_j with $j < i$ would have caused G_j, and hence G_i, to be reinitialized. For G_i to be the highest game reinitialized after a move, this means either that a player at level i moved within level i or to a lower level, or that a player moved to level i from a lower level. In both cases we can apply Lemma 6.2 to G_i (the first case is case (1) of the lemma, the second is case (2), and in both cases we take p_i to be the heaviest player who was used to define level i), which guarantees that we will have gained an improvement of more than $\frac{\varepsilon \gamma_i}{\alpha}$ in ϕ_i since the last initialization. □

With the above machinery in place, Theorem 6.1 follows from a straightforward induction:

Proof of Theorem 6.1: The proof proceeds by showing inductively that a game at level i can have at most $\lceil n_i \alpha \varepsilon^{-1} \log(n_i C) \rceil^{2^k - i}$ moves before reaching an ε-Nash equilibrium (or being prematurely ended by a higher game being reinitialized), where n_i is the number of players in G_i. The base case, $i = 2^k - 1$, is a game that is necessarily always at the bottom of the sequence. Any move that does not prematurely terminate this game must involve a player moving within this level, and is therefore by Lemma 6.3 a good move for G_i. Thus there can be at most $\lceil n_i \alpha \varepsilon^{-1} \log(n_i C) \rceil$ such moves before G_i reaches an ε-Nash equilibrium.

Now consider a game G_i for $i < 2^k - 1$; at any point it may either be the lowest level game in the sequence, or have a game G_{i+1} below it. In the first case, as above any move is either a good move for G_i or else ends G_i; in the second case, by induction G_{i+1} can run for at most $\lceil n_{i+1} \alpha \varepsilon^{-1} \log(n_{i+1} C) \rceil^{2^k - i - 1}$ steps before reaching an ε-Nash equilibrium. After this, the next move must necessarily be by a player outside G_{i+1}, i.e., by a player in G_i or a higher game, resulting in either a good move for G_i or termination of G_i. As G_i can only have $\lceil n_i \alpha \varepsilon^{-1} \log(n_i C) \rceil$ good moves before reaching its own ε-Nash equilibrium, and $n_{i+1} < n_i$, we obtain our result. □

Remarks:

1. In the special case where each strategy contains at most one exceptional edge, the exponent in Theorem 6.1 can be improved from 2^k to k. This is a natural case that arises when the edges are used to create "classes" of players, as discussed in the Introduction.

2. Theorem 6.1 holds for all variants of the dynamics considered here, including the unrestricted dynamics.

7 Open problems

We conclude by mentioning a few open problems arising from this work.

1. Can one extend the analysis of the ε-Nash dynamics to an arbitrary number of exceptional edges, and achieve at least subexponential convergence time? As mentioned earlier, this would actually cover all (not necessarily symmetric) congestion games. Also, can we say anything about the case of weighted players, where pure Nash equilibria may not necessarily exist?

2. Motivated by the polynomial time algorithm of [10] for finding an (exact) Nash equilibrium for symmetric *network* congestion games, can one show that the ε-Nash dynamics in such games converges rapidly (in an appropriate sense) to an *exact* equilibrium? Note that PLS-completeness is no longer an obstacle here.

3. What can one say about the properties of ε-Nash equilibria and their relationship to true equilibria? For example, it is not hard to see that the bounds on price of anarchy in [1,6] for exact equilibria with linear or polynomial delay functions carry over (with additional ε-dependent factors) to ε-equilibria, even for asymmetric games.

4. In our dynamics, only one player moves at each step. It would be interesting also to investigate dynamics in which all players move concurrently, as in [11] and [3] for the non-atomic case.

Acknowledgements

We thank Lisa Fleischer, Nicole Immorlica, Christos Papadimitriou, Tim Roughgarden and Kunal Talwar for helpful input.

References

[1] B. Awerbuch, Y. Azar and A. Epstein, "The price of routing unsplittable flow." *Proc. 37th ACM Symposium on Theory of Computing*, 2005, pp. 57–66.

[2] M. Beckmann, C.B. McGuire and C.B. Winsten, *Studies in the Economics of Transportation.* Yale University Press, 1956.

[3] A. Blum, E. Even-Dar and K. Ligett, "Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games." *Proc. 25th ACM Symposium on Principles of Distributed Computing*, 2006, pp. 45–52.

[4] X. Chen and X. Deng, "Settling the Complexity of 2-Player Nash-Equilibrium." *Proc. 47th IEEE Symposium on Foundations of Computer Science*, 2006, to appear.

[5] X. Chen, X. Deng and S.-H. Teng, "Computing Nash equilibria: Approximation and smoothed complexity." *Proc. 47th IEEE Symposium on Foundations of Computer Science*, 2006, to appear.

[6] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games." *Proc. 37th ACM Symposium on Theory of Computing*, 2005, pp. 67–73.

[7] C. Daskalakis, P.W. Goldberg and C.H. Papadimitriou, "The Complexity of Computing a Nash Equilibrium." *Proc. 38th ACM Symposium on Theory of Computing*, 2006, pp. 71–78.

[8] E. Even-Dar, A. Kesselman and Y. Mansour, "Convergence time to Nash equilibria." *Proc. 30th International Conference on Automata, Languages and Programming*, 2003, pp. 502–513.

[9] E. Even-Dar and Y. Mansour, "Fast convergence of selfish rerouting." *Proc. 16th ACM-SIAM Symposium on Discrete Algorithm*, 2005, pp. 772–781.

[10] A. Fabrikant, C.H. Papadimitriou and K. Talwar, "The complexity of pure Nash equilibria." *Proc. 36th ACM Symposium on Theory of Computing*, 2004, pp. 604–612.

[11] S. Fischer, H. Räcke and B. Vöcking, "Fast convergence to Wardrop equilibria by adaptive sampling methods." *Proc. 38th ACM Symposium on Theory of Computing*, 2006, pp. 653–662.

[12] M. Goemans, V. Mirrokni and A. Vetta, "Sink equilibria and convergence." *Proc. 46th IEEE Symposium on Foundations of Computer Science*, 2005, pp. 142–154.

[13] P.W. Goldberg, "Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game." *Proc. 25th ACM Symposium on Principles of Distributed Computing*, 2004, pp. 131–140.

[14] D.S. Johnson, C.H. Papadimitriou and M. Yannakakis, "How easy is local search?" *Journal of Computer and System Sciences* **37** (1988), pp. 79–100.

[15] M. Kearns and Y. Mansour, "Efficient Nash computation in large population games with bounded influence." *Proc. 18th Conference in Uncertainty in Artificial Intelligence*, 2002, pp. 259–266.

[16] R.J. Lipton, E. Markakis and A. Mehta, "Playing large games using simple strategies." *Proc. 4th ACM Conference on Electronic Commerce*, 2003, pp. 36–41.

[17] I. Milchtaich, "Congestion games with player-specific payoff functions," *Games and Economic Behavior* **13** (1996), pp. 111–124.

[18] V.S. Mirrokni and A. Vetta, "Convergence issues in competitive games." *Proc. 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2004, pp. 183–194.

[19] D. Monderer and L.S. Shapley, "Potential games." *Games and Economic Behavior* **14** (1996), pp. 124–143.

[20] J.F. Nash, "Equilibrium points in n-person games." *Proc. National Academy of Sciences* **36** (1950), pp. 48–49.

[21] J. Orlin, A. Punnen, and A. Schulz, "Approximate local search in combinatorial optimization." *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms*, 2004, pp. 587–596.

[22] C.H. Papadimitriou, "Algorithms, games and the Internet." *Proc. 33rd ACM Symposium on Theory of Computing*, 2001, pp. 749–753.

[23] R.W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria." *International Journal of Game Theory* **2** (1973), pp. 65–67.

[24] T. Roughgarden, *Selfish routing and the price of anarchy*. MIT Press, 2005.

[25] T. Roughgarden and E. Tardos, "Bounding the inefficiency of equilibria in nonatomic congestion games." *Games and Economic Behavior* **47** (2004), pp. 389–403.

[26] A.A. Schäffer and M. Yannakakis, "Simple local search problems that are hard to solve." *SIAM Journal on Computing* **20** (1991), pp. 56–87.

[27] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions." *Proc. 43rd IEEE Symposium on Foundations of Computer Science*, 2002, pp. 416–425.