Convergence Dynamics of Graphical Congestion Games

Richard Southwell[§], Yanjiao Chen[†], Jianwei Huang[§] and Qian Zhang[†]

[§]Information Engineering Department, The Chinese University of Hong Kong,[⋆]
Shatin, New Territories, Hong Kong
[†]Computer Science and Engineering Department, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong
{richardsouthwell254,chenyj.thu,jianweihuang}@gmail.com,qianzh@cse.ust.hk

Abstract. Graphical congestion games provide powerful models for a wide range of scenarios where spatially distributed individuals share resources. Understanding when graphical congestion game dynamics converge to pure Nash equilibria yields important engineering insights into when spatially distributed individuals can reach a stable resource allocation. In this paper, we study the convergence dynamics of graphical congestion games where players can use multiple resources simultaneously. We show that when the players are free to use any subset of resources the game always converges to a pure Nash equilibrium in polynomial time via lazy best response updates. When the collection of sets of resources available to each player is a matroid, we show that pure Nash equilibria may not exist in the most general case. However, if the resources are homogenous, the game can converge to a Nash equilibrium in polynomial time.

Key words: congestion game, resource allocation, matroid, games on graphs, graphical

1 Introduction

Congestion games have found applications in many scientific and engineering areas. The original congestion game model was introduced by Rosenthal [1]. The idea behind this model is that players select resources to use, and the payoff a player gains from using a given resource depends upon that resource's *congestion level* (i.e., the total number of players using it).

The original congestion game model is very general, because it allows different resources to be associated with different payoff functions, and it allows players to use multiple resources simultaneously. Also, the game has a very appealing feature called *the finite improvement property.*, which means that if the players

⋆ This work is supported by the General Research Funds (Project Number 412509) established under the University Grant Committee of the Hong Kong Special Administrative Region, China.

keep performing asynchronous better response updates (i.e., the players improve their strategy choices one at a time) then the system will eventually reach a pure Nash equilibrium - a strategy profile from which no player has any incentive to deviate. Intuitively, the finite improvement property means greedy updating always converges to a stable strategy profile.

The generality and pleasing convergence properties of the original congestion game model have led to its application to a wide range of resource allocation scenarios (e.g., economics [2], communication networks [3–6], network routing [7], network formation [8], ecology [9], and sociology [10]). However, the original model has the limitation that each player using the same resource gets the same payoff from it. Treating players identically is unsuitable for many scenarios in ecology [11], network routing [12], and wireless networks [13] where players are *heterogenous*. This has motivated many adaptations of the original congestion game, including congestion games with player-specific payoff functions [4,14] and weighted congestion games [16].

In [17], we considered the graphical congestion game (see Figure 1), an important generalization of the original congestion game concept. This model not only allows player-specific payoff functions but also models how the spatial positioning of the players affects their performance in the game. In the original congestion game model, any pair of users cause congestion to each another when using the same resource. In the graphical congestion game, we regard the players as vertices in a *conflict graph*. Only linked players cause congestion to each other. Unlinked players can use the same resource without causing congestion to each other. We describe some scenarios that can be modeled using graphical congestion games in Table 1.

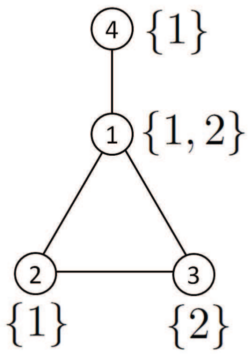

Fig. 1. A strategy profile in a graphical congestion game. The players (i.e., the vertices on the graph) select sets of resources to use. Player 1 is using resources 1 and 2. The amount of payoff a player gains from using a particular resource is a non-increasing function of the number of its neighbors who are also using that resource.

Although the graphical congestion game has a wide-range applications, it is no longer guaranteed to possess the finite improvement property or even pure

Table 1. How graphical congestion games can be used to model various resource sharing scenarios.

Scenario	Players Represent	Resources Represent	Links in the Conflict Graph Represent
Ecology [9]	Organisms	Food Sources or Habitats	Organisms are spatially close enough to compete for the same food source or habitat.
Wireless Networks [13, 15, 17]	Wireless Users	Channels	Users are close enough to cause significant interference to each other.
Market Sharing [23]	Businesses	Markets	Business locations are close enough to compete over the same customers.

Nash equilibria. Since the graphical congestion game is highly practically relevant yet may lose some nice features, the obvious question is as follows: *What are the conditions under which a graphical congestion game possesses a pure Nash equilibrium or even the finite improvement property?* This is a question of fundamental importance for many spatially distributed resource sharing scenarios.

1.1 Problem Definition

A generalized *graphical congestion game* is a 5-tuple $g = (\mathcal{N}, \mathcal{R}, (\zeta_n)_{n \in \mathcal{N}}, (f_n^r)_{n \in \mathcal{N}, r \in \mathcal{R}}, G)$, where:

- $\mathcal{N} = \{1, 2, ..., N\}$ is a set of N players.
- $\mathcal{R} = \{1, 2, ..., R\}$ is a set of R resources.
- $\zeta_n \subseteq 2^{\mathcal{R}}$ is[2] the collection of resource sets available to player $n \in \mathcal{N}$. During the game player n selects a member of ζ_n to use. Therefore ζ_n can be viewed as the set of strategies available to player n. Sometimes we refer to ζ_n as *the collection of available resource sets*, and the members of ζ_n as *available resource sets*.
- f_n^r is the non-increasing payoff function for a player $n \in \mathcal{N}$ using resource $r \in \mathcal{R}$.
- $G = (\mathcal{N}, \mathcal{E})$ is an undirected graph with vertex set $\mathcal{N}$ and edge set $\mathcal{E}$. Here $\mathcal{E}$ is a set of unordered pairs $\{n, n'\}$ of players. We say that player $n \in \mathcal{N}$ is linked to player $n' \in \mathcal{N}$ if and only if $\{n, n'\} \in \mathcal{E}$. We can interpret $\{n, n'\} \in \mathcal{E}$ as being equivalent to saying that n and n' can cause congestion to one another. We assume $\{n, n\} \notin \mathcal{E}$ for each player $n \in \mathcal{N}$. In other words, we assume that no player is adjacent to itself[3]. We refer to G as *the conflict graph*.

[2] Here $2^{\mathcal{R}}$ denotes the set of all subsets of $\mathcal{R}$.

[3] This is just a convention we adopt for simplicity. All our results persist if we allow players to be adjacent to/cause congestion to themselves, but the results would look more cumbersome. One could emulate the idea n is adjacent to themselves under our framework by replacing their payoff functions $f_n^r(x)$ with new payoff functions $f_n^r(x+1)$.

A *strategy profile* $\boldsymbol{X} \in \Pi_{n=1}^{N} \zeta_n$ consists strategy (i.e., a collection of resources) $X_n \in \zeta_n$ of each player $n \in \mathcal{N}$.

We define the *congestion level* $c_n^r(\boldsymbol{X})$ *of resource* $r \in \mathcal{R}$ *for player* $n \in \mathcal{N}$ *within strategy profile* $\boldsymbol{X}$ to be $c_n^r(\boldsymbol{X}) = |\{n' \in \mathcal{N} : \{n, n'\} \in \mathcal{E}, r \in X_{n'}\}|$. In other words $c_n^r(\boldsymbol{X})$ is the number of neighbors that n has in the conflict graph G which are using resource r in strategy profile $\boldsymbol{X}$. The *total payoff* that a player n gets in strategy profile $\boldsymbol{X}$ is

$$\sum_{r \in X_n} f_n^r\left(c_n^r(\boldsymbol{X})\right).$$

This is the sum of the payoffs $f_n^r\left(c_n^r(\boldsymbol{X})\right)$ that n receives from each of the resources r within the resource set X_n that n chooses.

1.2 Better, Best, and Lazy Best Response Updates

We are concerned with how graphical congestion games evolve through time as the players attempt to improve their resource choices.

Let us define an $[n] \to S$ *update* as the action where player $n \in \mathcal{N}$ switches to use resource set $S \in \zeta_n$, while all other players retain their existing resource selections. If the current strategy profile is $\boldsymbol{X}$, then the $[n] \to S$ update changes the strategy profile from $\boldsymbol{X}$ to a new strategy profile $\boldsymbol{Y} = (X_1, .., X_{p-1}, S, X_{p+1}, .., X_n)$. We wish to emphasize that an $[n] \to S$ update (and, in fact every update we consider) only involves *one* player changing its strategy, while all other players keep their strategy choices unchanged.

We say that an $[n] \to S$ update is a *better response update* if it improves player n's payoff, i.e.,

$$\sum_{r \in Y_n} f_n^r\left(c_n^r(\boldsymbol{Y})\right) > \sum_{r \in X_n} f_n^r\left(c_n^r(\boldsymbol{X})\right).$$

We say that $[n] \to S$ is a *best response update* if it improves player n's payoff to the maximum possible value among all better responses from the current strategy profile.

We say that $[n] \to S$ is a *lazy best response update* [16] if (a) $[n] \to S$ is a best response update, and (b) for any other best response update $[n] \to S'$ that n could perform, we have $|X_n - S'| + |S' - X_n| \geq |X_n - S| + |S - X_n|$. In other words, a lazy best response update is a best response update which minimizes the number $|X_n - S| + |S - X_n|$ of resources which n must add or remove from their currently chosen resource set X_n.

We say a strategy profile $\boldsymbol{X} \in \Pi_{n=1}^{N} \zeta_n$ is a *pure Nash equilibrium*[4] if and only if no better response updates can be performed by any player from $\boldsymbol{X}$.

We give an illustrative example of such a graphical congestion game in Figure 1. Suppose that the collections of available resources for the four players/vertices

[4] We always suppose players use pure strategies and so all of the Nash equilibria that we discuss are pure.

are $\zeta_1 = 2^{\{1,2,3\}}$, $\zeta_2 = \zeta_4 = \{\emptyset, \{1\}\}$, and $\zeta_3 = \{\emptyset, \{2\}\}$. Assume that the payoff functions are $f_n^r(x) = 1 - x$ for each player n and resource r. In the strategy profile $\boldsymbol{X}$ shown in Figure 1, player 1 uses strategy $X_1 = \{1, 2\}$ and receives a total payoff of $f_1^1(c_1^1(\boldsymbol{X})) + f_1^2(c_1^2(\boldsymbol{X})) = (1-2) + (1-1)$. From this strategy profile, player 1 could perform the better response update $[1] \to \{2\}$ (which is not a best response update), or the best response update $[1] \to \{2, 3\}$ (which is not a lazy best response update), or the lazy best response update $[1] \to \{3\}$ (which leads to a pure Nash equilibrium).

We are interested in how graphical congestion games evolve when the players keep performing better response updates. Nash equilibria are the fixed points of such dynamics, since no player has any incentive to deviate from a Nash equilibrium.

We can put properties a congestion game might possess in the ascending order of strength/desirability as follows;

1. A pure Nash equilibrium exists.
2. A sufficiently long sequence of lazy best response updates is guaranteed to drive the system to a pure Nash equilibrium.
3. A sufficiently long sequence of better response updates is guaranteed to drive the system to a pure Nash equilibrium (the finite improvement property).

This paper is mainly concerned with identifying conditions under which the generalized graphical congestion games have properties $1, 2$, and 3. It should be noted that the presence of property 3 implies the presence of property 2, which in turn, implies the presence of property 1. However it is possible to construct games with only subset (or none) of the above properties.

1.3 Previous Work

Graphical congestion games were introduced in [19], where the authors considered linear and non-player specific payoff functions. Such games are proved to have the finite improvement property when the graph is undirected or acyclic. But the authors illustrated a game on a directed graph with no pure Nash equilibria. In [20], players are assigned different weights, so they suffer more congestion from “heavier” neighbors. Both [19] and [20] restricted their attention to “singleton games” (where each player uses exactly one resource at any given time) with linear and non-player-specific payoff functions.

In [17], the authors introduced the more general graphical congestion game model as described in Section 1.1 to model spectrum sharing in wireless networks (see Table 1). The model allows generic player-specific payoff functions, as wireless users often have complicated and heterogeneous responses to the received interference. The authors showed that every singleton graphical congestion game with two resources has the finite improvement property. They also gave an example of a singleton graphical congestion game (with player-specific and resource-specific payoff functions) which does not possess any pure Nash equilibria. In [13], we extended upon this work by showing that every singleton graphical congestion game with *homogenous resources* (i.e. the payoff functions are not resource-specific) converges to a pure Nash equilibrium in polynomial time.

In [15], the authors investigated the existence of pure Nash equilibrium of spatial spectrum sharing games on general interference graphs, especially when Aloha and random backoff mechanisms are used for channel contention. They also proposed an efficient distributed spectrum sharing mechanism based on distributed learning.

1.4 Our results

We focus upon the generalized graphical congestion games where players can use multiple resources simultaneously. In general, a player n can use any available set of resources from ζ_n at any given time. Our results suggest that the kinds of restrictions put on the combinatorial structure of the collections of available resource sets, ζ_n, have a dramatic influence on whether the convergence properties exist or not.

In particular, we find that when the collections of available resource sets ζ_n are "matroids" [21], many powerful results can be derived. *A matroid* $M \subseteq 2^U$ *with a ground set* U is a set M of subsets $S \subseteq U$ (called independent sets) which has the following three properties[5];

1. Empty set $\emptyset \in M$.
2. If $S \in M$ and $S' \subseteq S$, then $S' \in M$.
3. If $S \in M$ contains less elements than $S' \in M$, then there exists some $x \in S' - S$ such that $S \cup \{x\} \in M$.

We refer to 1, 2 and 3 in the above list as the *matroid properties*. Properties 1 and 2 are natural. Property 3 ensures that many examples of "independent set structures" from combinatorics and linear algebra are matroids. In a graph, the collection of subsets of edges which hold no cycles is a matroid. If U is a finite set of vectors from a vector space, and M is the collection of linearly independent subsets of U, then M is a matroid. Another important example of a matroid is the *uniform matroid* $\{S \subseteq U : |S| \leq k\}$, which is the collection of all subsets of a set U which have no more than k elements. A simple kind of matroid is the *powerset* $M = 2^U = \{S \subseteq U\}$ (i.e., the collection of all subsets of U).

A *matroid graphical congestion game* is a graphical congestion game, within which the collection of available resource sets ζ_n of each player n is a matroid with a ground set $\mathcal{R}$. Matroids are very general, and so matroid graphical congestion games have many applications. In Table 1, we discussed how the graphical congestion game can be used to model ecologies, wireless networks, and market sharing. In each of these cases, it is more reasonable to assume that the collection of available resource sets of each player forms a uniform matroid than to treat the system as a singleton graphical congestion game. For example, in ecology

[5] We write $|S|$ to denote the number of elements in set S and $S' - S$ to denote the set of elements in S' but not S.

the organisms will be able to access multiple food sources, but they will not be able to access more than a certain number of food sources because of limited time and energy. In wireless networks, users can divide their transmission power among many channels, however they cannot access too many channels because their total power is limited.[6] In market sharing games (e.g., [23]), each player has a fixed budget they can spend upon serving markets. When the cost of serving each market is the same, this corresponds to a uniform matroid congestion game because the number of markets a player can serve is capped. Linked payers in a matroid graphical congestion game could represent businesses who are close enough to compete for the same customers. As [16] noted, some network formation games correspond to congestion games with a matroid structure. For example, in [22] the authors considered the game where players select spanning trees of a graph, but suffer congestion from sharing edges with other players. In such scenarios, the conflict graph could represent which players are able to observe each others' actions.

In section 2, we consider the properties of a special important type of matroid graphical congestion game, the *powerset graphical congestion game*, within which the collection of available resource sets ζ_n of each player n is a powerset $\zeta_n = 2^{Q_n}$ for some subset $Q_n \subseteq \mathcal{R}$ of available resources. In section 3, we investigate the properties of more general matroid graphical congestion games. Our main results are listed below (and illustrated in Figure 2);

- There exist powerset graphical congestion games with homogenous resources which do not have the finite improvement property (Theorem 1)
- Every powerset graphical congestion game will reach a pure Nash equilibrium in polynomial time when the system evolves via lazy best response updates (Theorem 2).
- There exist matroid graphical congestion games which posses no pure Nash equilibria (Theorem 3).
- Every matroid graphical congestion game with homogenous resources will reach a pure Nash equilibrium in polynomial time when the system evolves via lazy best response updates (Theorem 4).

Our main result is Theorem 4, because it identifies a very general class of games with pleasing convergence properties. This result is especially meaningful for wireless networks, because wireless channels often have equal bandwidth, which means that they correspond to homogenous resources (under flat fading or interleaved channelization). The way we prove this convergence result is to define a *potential function*, which decreases whenever a player performs a lazy best response update. The existence of such a function guarantees that lazy best response updatings eventually lead to a fixed point (a pure Nash equilibrium).

Due to limited space, we refer the readers to our online technical report [24] for the full proofs of most results in this paper.

[6] In reality, when a user shares its power among many channels, the benefit they receive from using each one is diminished. Our game model does not capture this effect, however other models that do [18] are often analytically intractable.

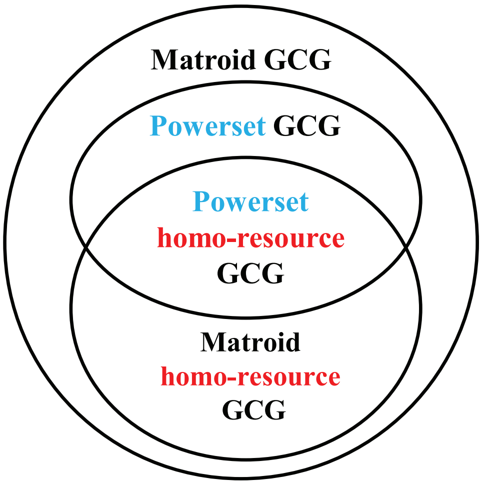

Fig. 2. In both **Powerset GCG** and **Matroid homo-resource GCG**, lazy best response update converges to pure Nash equilibria in polynomial time. However, even in the intersection class of **Powerset homo-resource GCG**, there exist examples where better response update may never converge to pure Nash equilibria.

2 Powerset Graphical Congestion Games

We begin our exploration of the dynamics of graphical congestion games with the "powerset" case, where players may use any subset of a set Q_n of resources available to them. In powerset congestion games, the decision of whether or not to use one resource has no effect on a player's ability to use the other resources. This fact allows us to decouple the system and consider the usage each resource separately.

As we shall see, the players in a powerset graphical congestion game can reach a pure Nash equilibrium in polynomial time via selfish updating. However, the players must be careful about what kind of updates they perform, because the following result suggests that better response updating is not guaranteed to lead to a pure Nash equilibrium.

Theorem 1. *There exist powerset graphical congestion games with homogenous resources which do not have the finite improvement property.*

Proof. Consider the powerset graphical congestion game g with players $\mathcal{N} = \{1, 2, 3\}$, resources $\mathcal{R} = \{1, 2, 3, 4\}$, strategy sets $\zeta_1 = \zeta_2 = \zeta_3 = 2^{\{1,2,3,4\}}$ and payoff functions f_n^r such that $(f_1^r(0), f_1^r(1), f_1^r(2)) = (0, -5, -7)$ and $(f_2^r(0), f_2^r(1), f_2^r(2)) = (f_3^r(0), f_3^r(1), f_3^r(2)) = (0, -2, -7)$. The game is is played on a three vertex complete graph G. Figure 3 shows how better response updating can lead to cycles in g, meaning g does not have the finite improvement property. □

Notice that the example game in the proof of Theorem 1 is played on a complete graph and has homogenous resources. Thus the lack of finite improvement

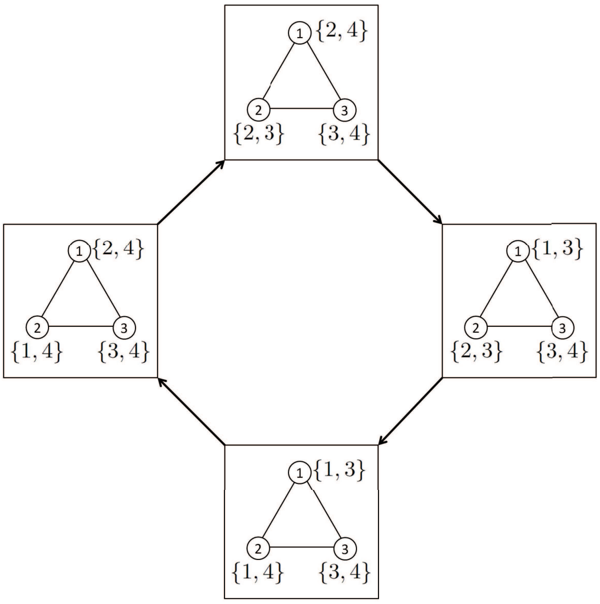

Fig. 3. A cycle in the best response dynamics of the powerset graphical congestion game discussed in the proof of Theorem 1. The arrows represent how the strategy profile changes with certain better response updates. Better response updating cannot be guaranteed to drive this game into a pure Nash equilibrium because better response updating can lead to cycles.

property is not due to either special property of the graph or the resources. Theorem 1 seems to be quite negative. However, as we shall see, the players often *can* be gaurenteed to reach pure Nash equilibria if they update their resources in special ways (instead of unregulated asynchronous updates). Before we describe this in more details, let us introduce some tools that will be useful throughout our analysis: beneficial pickups, beneficial drops, and the temperature function.

2.1 Beneficial Pickups and Drops

A better response update may alter the set of resources that a player is using in quite complicated ways. However, we will show that better response updates can be decomposed into sequences of elementary update operations. Here we introduce two such operations: the beneficial pickup (where a player starts using a good new resource) and the beneficial drop (where a player stops using a bad old resource).

More formally, suppose we have a graphical congestion game in the strategy profile $\boldsymbol{X}$. *A beneficial pickup* is a better response update $[n] \to X_n \cup \{a\}$ with $a \notin X_n$ (i.e., a beneficial pickup is where a player starts using a new resource a and obtains additional benefits). A *beneficial drop* is a better response update

$[n] \to X_n - \{b\}$ where $b \in X_n$ (i.e., a beneficial drop is where a player stops using a resource b and gains benefits).

To illustrate these concepts, consider the graphical congestion game depicted in Figure 1 with parameters as described in Section 1.2. In this case, $[1] \to \{1, 2, 3\}$ is a beneficial pickup that player 1 can perform and $[1] \to \{2\}$ is a beneficial drop that player 1 can perform.

We can use beneficial pickups and drops to construct more complex updates. Thinking in this way is useful, because we can define a global "temperature" function which decreases every time a player conducts a beneficial pickup or drop.

2.2 The Temperature Function

The temperature function maps strategy profiles to integers. In certain scenarios, the temperature function acts like a potential function, which decreases with lazy best response updates[7]. This fact allows us to prove our polynomial time convergence results.

To build the temperature function, we associate each payoff function f with a left-threshold value $T_N^{\leftarrow}[f]$ (which, roughly speaking, is the maximum integer x such that $f(x) \geq 0$) and a right-threshold value $T_N^{\rightarrow}[f]$ (which, roughly speaking, is the minimum integer x such that $f(x) \leq 0$). The values of these thresholds also depend on the integer N. We will take N to be the number of players in our game when we apply these concepts later.

More precisely, suppose f is a non-increasing function and N is an integer. We define *the left-threshold* $T_N^{\leftarrow}[f]$ *of* f *with respect to* N as follows:

$$T_N^{\leftarrow}[f] = \begin{cases} -1, & \text{if } f(x) < 0, \forall x \in \{0, ..., N-1\}, \\ \max\{x \in \{0, .., N-1\} : f(x) \geq 0\}, & \text{otherwise.} \end{cases}$$

We define *the right-threshold* $T_N^{\rightarrow}[f]$ *of* f *with respect to* N as follows:

$$T_N^{\rightarrow}[f] = \begin{cases} N, & \text{if } f(x) > 0, \forall x \in \{0, ..., N-1\}, \\ \min\{x \in \{0, .., N-1\} : f(x) \leq 0\}, & \text{otherwise.} \end{cases}$$

In an N-player graphical congestion game the input of a payoff function f will be a congestion level in the range $\{0, 1, ..., N-1\}$. The following lemma describes how $T_N^{\leftarrow}[f]$ and $T_N^{\rightarrow}[f]$ indicate when a resource's congestion level is so high that it is no longer worth using.

Lemma 1. *Suppose* $T_N^{\leftarrow}[f]$ *and* $T_N^{\rightarrow}[f]$ *are the left-threshold and right-threshold values of the non-increasing function* f *(with respect to* N*), then for any* $x \in \{0, ..., N-1\}$*,*

- $f(x) > 0$ *if and only if* $x \leq T_N^{\rightarrow}[f] - 1$*, and*

[7] The temperature function is not always a potential function, because it may not decrease when certain better response updates are performed in certain cases.

- $f(x) < 0$ *if and only if* $x \geq T_N^{\leftarrow}[f] + 1$.

Lemma 1 can be proved using basic facts about non-increasing functions. With this lemma in place we shall define the temperature function.

The temperature function Θ *associated with an* N*-player graphical congestion game* g is defined as

$$\Theta(\boldsymbol{X}) = \sum_{n \in \mathcal{N}} \sum_{r \in X_n} \left(c_n^r(\boldsymbol{X}) - T_N^{\leftarrow}[f_n^r] - T_N^{\rightarrow}[f_n^r] \right).$$

In many types of graphical congestion game, the temperature function always decreases with lazy best response updates. Now we will show that the temperature function decreases every time a player performs a beneficial pickup or drop.

Lemma 2. *Suppose that we have a graphical congestion game in a strategy profile* $\boldsymbol{X}$*, and a player* n *performs a beneficial pickup,* $[n] \rightarrow X_n \cup \{a\}$*, which drives the system into a strategy profile* $\boldsymbol{Y}$*. We have* $\Theta(\boldsymbol{Y}) \leq \Theta(\boldsymbol{X}) - 1$.

Lemma 2 can be proved using Lemma 1 together with the fact that $f_n^a(c_n^a(\boldsymbol{X})) > 0$ whenever $[n] \rightarrow X_n \cup \{a\}$ is a beneficial pickup.

Lemma 3. *Suppose that we have a graphical congestion game in a strategy profile* $\boldsymbol{X}$*, and a player* n *performs a beneficial drop,* $[n] \rightarrow X_n - \{b\}$*, which drives the system into a strategy profile* $\boldsymbol{Y}$*. We have* $\Theta(\boldsymbol{Y}) \leq \Theta(\boldsymbol{X}) - 1$.

Lemma 3 can be proved using Lemma 1 together with the fact that $f_n^b(c_n^b(\boldsymbol{X})) < 0$ whenever $[n] \rightarrow X_n - \{b\}$ is a beneficial drop.

The temperature function clearly takes integer values. Another crucial feature of the temperature function is that it is bounded both above and below.

Lemma 4. *If* $\boldsymbol{X}$ *is a strategy profile of a graphical congestion game with* N *players and* R *resources, then temperature function* Θ *satisfies the inequalities* $R(N - 2N^2) \leq \Theta(\boldsymbol{X}) \leq RN^2$.

2.3 Convergence Dynamics of Powerset Graphical Congestion Games

Lemma 5 characterizes the relationship between the lazy best response and the beneficial pickups and drops.

Lemma 5. *In a powerset graphical congestion game, every lazy best response can be decomposed into a sequence of beneficial pickups and/or beneficial drops.*

We know from Lemmas 2 and 3 that beneficial pickups and drops decreases the temperature function. Hence Lemma 5 essentially shows that the temperature function is a potential function, which decreases by integer steps when a powerset graphical congestion game evolves via lazy best response updates.

Theorem 2. *Consider a powerset graphical congestion game with N players and R resources. A Nash equilibrium can be reached from any initial strategy profile within $R(3N^2 - N)$ asynchronous lazy best response updates.*

Sketch of proof Since each beneficial pickup or drop decreases the temperature function Θ by at least one (Lemmas 2 and 3), and each lazy best response update can be decomposed into beneficial pickups and drops (Lemma 5), we have that each lazy best response update decreases the temperature function by at least one. Since the temperature function is bounded above by RN^2 and below by $R(N-2N^2)$ (Lemma 4), then no more than $RN^2-(R(N-2N^2)) = R(3N^2-N)$ lazy best response updates can be performed starting from any strategy profile. When no more lazy best response update can be performed, we reach a pure Nash equilibrium. □

3 Matroid Graphical Congestion Games

Powerset graphical congestion games have a relatively simple combinatorial structure, which allows us to prove with relative ease that they always have pure Nash equilibria. When the resource availability sets ζ_n's have a more complicated structure, this is no longer true. In this section, we shall investigate the properties of the more general *matroid graphical congestion games*, where each player's collection of available resource sets ζ_n is a matroid. We start by showing that in a pure strategy Nash equilibrium may not exist in general.

Theorem 3. *There exist matroid graphical congestion games which do not possess a pure Nash equilibrium.*

Sketch of proof In [17], the authors gave an example of a singleton graphical congestion game g (with strictly positive payoff functions) that has no pure Nash equilibria. We can convert g into a matroid graphical congestion game g' by giving players the extra option of using no resources (i.e., by adding the empty set into their collection of available resource sets). Since using a resource in g' leads to a positive payoff, rational players in g' will behave exactly as in g (i.e., they will always want to use some resource). Since g has no pure Nash equilibria, g' has no pure Nash equilibria either. □

Next we shall examine a special type of matroid graphical congestion game, which is guaranteed to possess a pure Nash equilibrium and nice convergence properties.

3.1 Convergence dynamics of matroid graphical congestion games with homogenous resources

We say a graphical congestion game $g = (\mathcal{N}, \mathcal{R}, (\zeta_n)_{n\in\mathcal{N}}, (f_n^r)_{n\in\mathcal{N},r\in\mathcal{R}}, G)$ *has homogenous resources* when the payoff functions are not resource specific (i.e.,

$f_n^1(x) = f_n^2(x) = ... = f_n^R(x) = f_n(x), \forall n \in \mathcal{N}, \forall x)$. Note that different players can have different payoff functions. When discussing resource homogenous games, we often suppress the superscript on the payoff functions, writing $f_n^r(x)$ as $f_n(x)$ to represent the fact that the payoff functions do not depend on the resources.

We will show that a matroid graphical congestion game with homogenous resources will reach a pure Nash equilibrium in polynomial time if the players perform lazy best response updates. We prove this result with the help of the temperature function. Before we do this, we must introduce a third type of elementary update operation – the beneficial swap, which is a better response update $[n] \rightarrow (X_n \cup \{a\}) - \{b\})$ where $a \notin X_n$ and $b \in X_n$ (i.e., a beneficial swap is where a player stops using a resource b and starts using a resource a, and benefits as a result.)

Our next result states that in any graphical congestion game with homogenous resources (but not necessarily with matroid structure), a beneficial swap will decrease the temperature function Θ by at least one.

Lemma 6. *Suppose we have a graphical congestion game with homogenous resources in a strategy profile* $\boldsymbol{X}$*, and we perform a beneficial swap* $[n] \rightarrow (X_n \cup \{a\}) - \{b\}$*, which moves the system into a strategy profile* $\boldsymbol{Y}$*. We have* $\Theta(\boldsymbol{Y}) \leq \Theta(\boldsymbol{X}) - 1$*.*

Lemma 6 follows from the fact that if $[n] \rightarrow (X_n \cup \{a\}) - \{b\}$ is a beneficial swap and the resources are homogenous, then $c_n^a(\boldsymbol{X}) < c_n^b(\boldsymbol{X})$.

Lemmas 2, 3, and 6 together imply that any beneficial pickup, drop, or swap in a graphical congestion game with homogenous resources will decrease the temperature function. Next we will show that if the strategy sets ζ_n's of the game are *matroids*, then it is always possible to perform a beneficial pickup, drop, or swap from a non-equilibrium state. In particular we will show that each lazy best response update in a matroid graphical congestion game with homogenous resources can be decomposed into a sequence of beneficial pickups, drops, and/or swaps. The following three lemmas will allow us to achieve this goal.

Lemma 7. *If* $[n] \rightarrow S$ *is a lazy best response update that can be performed from a strategy profile* $\boldsymbol{X}$ *of a matroid graphical congestion game with homogenous resources and* $|X_n| < |S|$*, then there exists* $a \in S - X_n$ *such that* $[n] \rightarrow X_n \cup \{a\}$ *is a beneficial pickup that player* n *can perform from* $\boldsymbol{X}$*.*

Lemma 8. *If* $[n] \rightarrow S$ *is a lazy best response update that can be performed from a strategy profile* $\boldsymbol{X}$ *of a matroid graphical congestion game with homogenous resources and* $|X_n| > |S|$*, then there exist* $b \in S - X_n$ *such that* $[n] \rightarrow X_n - \{b\}$ *is a beneficial drop that player* n *can perform from* $\boldsymbol{X}$*.*

Lemma 9. *If* $[n] \rightarrow S$ *is a lazy best response update that can be performed from a strategy profile* $\boldsymbol{X}$ *of a matroid graphical congestion game with homogenous resources and* $|X_n| = |S|$*, then there exists* $a \in X_n - S$ *and* $\exists b \in S - X_n$ *such* *that* $[n] \to (X_n \cup \{a\}) - \{b\}$ *is a beneficial swap that player* n *can perform from* $\boldsymbol{X}$.

Lemmas 7 and 8 can be shown using the basic matroid properties. Our proof to Lemma 9 uses a more sophisticated result about matroids shown in [21]. With Lemmas 7, 8, and 9, we can prove the following main result of this paper.

Theorem 4. *Consider a matroid graphical congestion game with homogenous resources with* N *players and* R *resources. A Nash equilibrium can be reached from any initial strategy profile within* $R(3N^2 - N)$ *asynchronous lazy best response updates.*

Sketch of proof Since each beneficial pickup, drop, or swap decreases the temperature function Θ by at least one (Lemmas 2, 3, and 6) and each lazy best response update can be decomposed into beneficial pickups, drops, or swaps (as can be proved inductively using Lemmas 7, 8, and 9), we have that each lazy best response update decreases the temperature function by at least one. Since the temperature function is bounded above by RN^2 and below by $R(N - 2N^2)$ (Lemma 4), no more than $RN^2 - (R(N - 2N^2)) = R(3N^2 - N)$ lazy best response updates can be performed starting from any strategy profile. When no more lazy best response update can be performed, we reach a pure Nash equilibrium. □

By considering Theorem 4 in conjunction with Theorem 1, we can see an interesting separation between the dynamics that always reach a pure Nash equilibrium and the dynamics which sometimes do not. Theorem 1 implies the existence of matroid graphical congestion games with homogenous resources that will never converge to pure Nash equilibria when the players do better response updates. However, Theorem 4 implies that when the players restrict themselves to lazy best response updates (which are more accurate and rational), they *are* guaranteed to reach a pure Nash equilibrium in polynomial time.

4 Conclusion

We have derived many results which are useful for understanding when graphical congestion games converge to pure Nash equilibria. Theorem 1 is quite negative, because it implies the existence of games with simple features (players that can use any combination of resources, and resources are homogenous) which cannot be guaranteed to converge to pure Nash equilibria under generic better response updating. However, Theorems 2 and 4 imply that in many cases (powerset games, or matroid games with homogenous resources) the players do converge to pure Nash equilibria under lazy best response updating. These results are very encouraging, because they imply that spatially distributed individuals will quickly be able to organize themselves into a pure Nash equilibrium in a wide range of scenarios. Just so long as the players are rational enough to restrict themselves to lazy best response updates. We obtained our convergence results by breaking better response updates into more elementary operations, and observing how these operations alter the value of the temperature function we defined. In the future, we will use these results to study the convergence dynamics of more general games, where players have generic collections of available resource sets.

References

1. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory. 2, 65–67, (1973)
2. Tennenholtz, M., Zohar, A.: Learning equilibria in repeated congestion games. Proceedings of AAMAS ’09. (2009)
3. Liu, M., Wu, Y.: Spectum sharing as congestion games. In: Proceedings the 46th Annual Allterton Conference on Communication, Control, and Computing. (2008)
4. Law, L., Huang, J., Liu, M., Li, S.: Price of Anarchy for Cognitive MAC Games. In: Proceedings of IEEE GLOBECOM. (2009)
5. Chen, X., Huang, J.: Evolutionarily Stable Spectrum Access in a Many-Users Regime. In: Proceedings of IEEE GLOBECOM. (2011)
6. Southwell, R., Huang, J.: Spectrum Mobility Games. In: IEEE INFOCOM. (2012)
7. Vöcking, B., Aachen, R.: Congestion games: Optimization in competition. In: Proceedings of the Second ACiD Workshop. (2006)
8. Tardos, E., Wexler, T.: Network formation games and the potential function method. In: Algorithmic Game Theory, chapter 19. 487–516, (2007)
9. Fretwell, S. D., Lucas, H. L.: On Territorial Behavior and Other Factors Influencing Habitat Distribution in Birds. In: Acta Biotheor. 19, 16–36, (1969)
10. Lachapelle, A., Wolfram, M.: On a mean field game approach modeling congestion and aversion in pedestrian crowds. In: Transportation Research Part B: Methodological. 45, 1572-1589, (2011)
11. Godin, J., Keenleyside, M.: Foraging on Patchily Distributed Preyby a Cichlid Fish (Teleostei, Cichlidae): A Test of the Ideal Free Distribution Theory. In: Anim. Behav. 32, 120-131, (1984)
12. Ackermann, H., Röglin, H., Vöcking, B.: On the Impact of Combinatorial Structure on Congestion Games. In: Proceedings of FOCS ’06. (2006)
13. Southwell, R., Huang, J.: Convergence dynamics of resource-homogeneous congestion games. In: Proceedings of GameNets ’11. (2011)
14. Milchtaich, I.: Congestion Games with Player-Specific Payoff Functions. In: Games and Economic Behavior. 13, 111–124, (1996)
15. X. Chen and J. Huang: Spatial Spectrum Access Game: Nash Equilibria and Distributed Learning, In: ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), Hilton Head Island, South Carolina, USA, June 2012
16. Ackermann, H., Röglin, H., Vöcking, B.: Pure Nash Equilibria in Player-Specific and Weighted Congestion Games. In: Internet and Network Economics, LNiCS. 4286, 50–61, (2006)
17. Tekin, C., Liu, M., Southwell, R., Huang, J. Ahmad, S.: Atomic Congestion Games on Graphs and its Applications in Networking. To Appear In: IEEE Transactions on Networking. (2012)
18. Etkin, R., Parekh, A., Tse, D.: Spectrum sharing for unlicensed bands. In: IEEE Journal on Selected Areas in Communications. 25, 517–528, (2007)
19. Bilo, V., Fanelli, A., Flammini, M., Moscardelli, L.: Graphical congestion games. In: Algorithmica. 61, 274–297, (2008)

20. Fotakis, D., Gkatzelis, V., Kaporis, A., Spirakis, P.: The impact of social ignorance on weighted congestion games. In: Internet and Network Economics, LNiCS. 5929, 316–327, (2009)
21. Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency. In: Matroids, Trees, Stable Sets, Volume B. 39–69, (2009)
22. Werneck, R, Setubal, J., Conceicao, A.: Finding minimum congestion spanning trees. In: Journal of Experimental Algorithmics. 5, (2000)
23. Goemans, M., Li, L., Mirrokni, V., Thottan, M.: Market sharing games applied to content distribution in ad-hoc networks. In: Proceedings of MobiHoc '04, (2004)
24. Southwell, R., Chen, Y., Huang, J., Zhang, Q.: Convergence Dynamics of Graphical Congestion Games, Technical Report, http://jianwei.ie.cuhk.edu.hk/publication/GCCConvergenceTechReport.pdf