Fast and Compact: A Simple Class Of Congestion Games

Samuel Ieong Robert McGrew Eugene Nudelman Yoav Shoham Qixiang Sun
Computer Science Department
Stanford University, Stanford CA 94305
Email: {sieong,bmcgrew,eugnud,shoham,qsun}@cs.stanford.edu

Abstract

We study a simple, yet rich subclass of congestion games that we call *singleton* games. These games are exponentially more compact than general congestion games. In contrast with some other compact subclasses, we show tractability of many natural game-theoretic questions, such as finding a sample or optimal Nash equilibrium. For best- and better-response dynamics, we establish polynomial upper and lower bounds on the rate of convergence and present experimental results. We also consider a natural generalization of singleton games and show that many tractability results carry over.

Introduction

Game theory is the central tool used in artificial intelligence to understand multiagent systems with strategic agents. *Congestion games* (Rosenthal 1973) have been an active area of research in computer science because they can model diverse phenomena such as processor scheduling, routing, and network design. In these games each agent (e.g., a task owner) is allowed to choose a subset of a global set of resources (e.g., CPUs), and agents' costs depend only on the number, but not the identities, of the other agents using the same resources. A desirable outcome of multiagent interactions is that agents coordinate on an action profile which is stable against strategic deviation — a *Nash equilibrium* (NE). Unlike general games, congestion games always have pure-strategy equilibria. Moreover, NEs in congestion games can be justified even for boundedly-rational agents, since *better-response* and *best-response* dynamics — very natural asynchronous dynamics in multiagent systems — are both guaranteed to converge to a NE in a finite number of steps. Unfortunately, most natural problems in congestion games have been shown to be intractable: even representing general congestion games requires space exponential in natural parameters of the game. Similarly, the problem of finding Nash equilibria is suspected to take time also exponential in natural parameters, as is the convergence of best response dynamics.

However, little attention has yet been paid to studying interesting classes of congestion games that are computationally tractable. In this paper, we focus on *singleton congestion games*: a tractable class of congestion games in which each player is restricted to choosing a single resource. With a careful choice of cost functions, singleton games can capture many settings naturally modeled by general congestion games. For example, consider a group of factories in a region, each of whom has a choice of suppliers for a given input. Note that, in contrast to much previous work on congestion games, we make no simplifying assumptions on the cost functions. In our example, the suppliers may face economies of scale at one level of production, but diseconomies of scale at another. Thus, simplifying restrictions on cost functions, such as assuming monotonicity, are insufficient for modeling the full range of possible phenomena.

In this paper, we show the following properties of singleton games:

- They are representable in polynomial space.
- Even optimal equilibria can be found in time polynomial in the size of the representation.
- Both best- and better-response dynamics are guaranteed to converge to an equilibrium in polynomial time.

In addition to being interesting and expressive in their own right, singleton games allow a reduction from a broader class of games we call Independent Resource Congestion Games (IRCGs), in which players may choose multiple resources and may be asymmetric in their choices. These games can also be represented in polynomial space, and the reduction shows that best-response dynamics is guaranteed to converge in polynomial time in IRCGs as well.

Related Work

Much recent effort in computer science has gone into the extremely challenging (Papadimitriou 2001) problem of finding a Nash equilibrium for general games (Conitzer & Sandholm 2003; Porter, Nudelman, & Shoham 2004). Given the generality of the problem, however, intractability results are common while tractability results are rare. Equilibria tend not to be unique and are computationally hard to find. Also, best-response dynamics is not guaranteed to converge. As a result, in many games it is hard to justify a Nash equilibrium, especially a mixed-strategy one, as a meaningful concept. Moreover, capturing strategic interactions in normal form requires space exponential in the number of players, and, hence, is hard even without computational issues.

A complementary research direction has been to introduce compact game representations (Koller & Milch 2001; Kearns, Littman, & Singh 2001; Leyton-Brown & Tennenholtz 2003; Bhat & Leyton-Brown 2004). Most of these representations strive to be as general as possible while capturing a particular kind of independence. However, a form that can represent any game must necessarily allow games which do not have pure-strategy Nash equilibria or convergent best-response dynamics. Similarly, this line of research rarely results in polynomial time algorithms.

One solution to this dilemma has been to study particular classes of games which have nice properties or reflect computer science practice. Congestion games and subclasses have been targeted by this line of research.

Fabrikant *et al.* 2004, looked at a very natural subclass called network congestion games, in which players route flow through a network. Unfortunately, they showed that finding a Nash equilibrium in general network congestion games is complete for the complexity class PLS (a semantic class generalizing local search problems), by a reduction which also showed that best-response dynamics may have convergence time exponential in the input size.

Other work (Even-Dar, Kesselman, & Mansour 2003; Anshelevich *et al.* 2004) considers the convergence time of best response in congestion games with a particular cost function and resource set. We believe that our work is the first to allow arbitrary cost functions and make assumptions only about the choice of resources.

Finally, singleton games, sometimes called simple games, have been studied in a variety of contexts in economics (Holzman & Law-Yone 1997), but this work has usually been non-computational. Milchtaich 1996 considers computational issues in a generalization of singleton games, but his discussion is heavily dependent on monotonic cost assumptions. The existence of pure-strategy Nash equilibria in other generalizations of singleton games, usually assuming some variant of monotonic costs, has been extensively studied as well; see Voorneveld *et al.* 1999 for a survey.

Notation and Background

A *game in normal form* (for an introduction to game theory see, e.g. (Osborne & Rubinstein 1994)) is a tuple $\langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$, where $N = \{1, \ldots, n\}$ is a finite set of players, S_i is a finite set of actions available to player i. Let $S = \prod_{i\in N} S_i$. Then $u_i : S \to \mathbb{R}$ is the utility function for each player i that maps a profile of actions to a value. Let $s \in S$ denote a joint strategy profile of players. We will use s_{-i} to stand for the same profile with i's strategy excluded, so that (s_i, s_{-i}) forms a complete profile of actions.

Define the *best response of player i to s_{-i}* to be the set of actions yielding optimal utility given s_{-i}: $BR_i(s_{-i}) = \text{argmax}_{s_i\in S_i} u_i(s_i, s_{-i})$. Similarly, the *better response of i to s* is $\Delta R_i(s) = \{s'_i \in S_i : u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})\}$.

An action profile s is a *(pure strategy) Nash equilibrium* (NE) if every player is playing best response to his opponents. Formally: for all i, $s_i \in BR_i(s_{-i})$ or, equivalently, for all i, $\Delta R_i(s) = \emptyset$. In this paper we will only consider pure strategy equilibria.

Best-response and better-response dynamics refer to the strategy profiles over a series of best-response or better-response moves. We say that best-response (better-response) dynamics *converges in m steps* if any series of m best-responses (better-responses) yields a NE.

Congestion Games

A *congestion game* (CG) is a tuple $G = \langle N, R, (S_i)_{i\in N}, (c_j)_{j\in R}\rangle$, where N is the set of players with size n, R is the set of resources with size k, $S_i \subseteq 2^R$ is the action space of player i, and $c_r : \{1, 2, \ldots, n\} \to \mathbb{R}$ is the cost function for resource r.

Denote the number of players choosing a particular resource r by $\nu_r(s)$. We will refer to vectors $\nu = \nu(s)$ as *configurations* of players in G. The semantics of CGs is that each player i choosing a set s_i of resources must pay the total cost $c^{(i)}(s) = \sum_{r\in s_i} c_r(\nu_r(s))$. We will assume that players are attempting to minimize cost. Thus, each congestion game defines a normal-form game $\langle N, (S_i), (-c^{(i)})\rangle$.

To represent a general congestion game, we must define two things. First, the cost function for all resources can be represented as a *cost matrix* $C = (c_i(j))_{ij}$ of size $k \times n$. This can be done in space $O(nk)$ even for arbitrary cost functions. Second, the set S_i for each player must be represented. This requires space exponential in k for general congestion games.

Note that, for a fixed number of players, the CG representation is only polynomially smaller than normal form. Thus, a NE can be found in polynomial time by expanding the CG representation into a normal form game. However, this can take exponential time for games which have compact representations or unbounded numbers of players.

A concept that is very tightly related to congestion games is that of a potential function. For a normal form game G, a function $\Phi : S \to \mathbb{R}$ is called an *(exact) potential function*, if for all s, s' such that $\exists i \;\; s_{-i} = s'_{-i}$, we have $\Phi(s) - \Phi(s') = u_i(s') - u_i(s)$.[1] It is easy to see that any profile which minimizes the potential function is a Nash equilibrium, since each best- or better-response move must decrease the potential. Thus, games with potential functions always have pure-strategy NEs. Rosenthal (1973) showed that CGs always have exact potential functions, given by $\Phi(s) = \sum_{r\in R} \sum_{l=1}^{\nu_r(s)} c_r(l)$.

The *social cost* of a configuration ν for a cost matrix C, denoted by $K(\nu, C)$, is the sum of the costs paid by all players: $K(\nu, C) = \sum_r \nu_r \cdot C(r, \nu_r)$. A configuration ν is a *socially optimal Nash equilibrium* if ν is a Nash equilibrium and for all other NE ν', $K(\nu, C) \leq K(\nu', C)$.

Singleton Congestion Games

Since the representation size for general congestion games may be exponential, it is natural to consider compactly representable subclasses. We first focus on singleton congestion games, where each player is allowed to choose any resource

[1] Since our games use costs rather than utilities, we adopt the convention that potential decreases with a decrease in cost or an increase in utility.

from R, but must choose exactly one. As described in the introduction, this class is rich enough to model many interesting interactions.

Definition 1. *A* singleton congestion game *(SCG) is a congestion game* $\langle N, R, (S_i)_{i\in N}, (c_j)_{j\in R}\rangle$ *with* $S_i = \{X \in 2^R : |X| = 1\}$.

Notice that an SCG is completely specified in polynomial space by its $k \times n$ cost matrix. We also note that players are anonymous in SCGs. So, given a configuration ν, all strategy profiles s with $\nu = \nu(s)$ differ only by a permutation of players. Thus, we can refer to equilibrium *configurations*, instead of strategy profiles.

Despite the simplicity of SCGs, with an arbitrary cost matrix C there may be exponentially many Nash equilibria, as the following cost matrix for $2n$ players and k resources demonstrates:

$$C(i,j) = \begin{cases} 0 & \text{if } j \text{ is even} \\ \infty & \text{if } j \text{ is odd} \end{cases}$$

A configuration is a Nash equilibrium for C if there are an even number of players at each resource. In this game, there are $\binom{k+n-1}{k-1}$ different possible NE configurations, which is exponential even without considering the permutation of players. It is also easy to construct SCGs with very few NEs.

Given these observations and the fact that SCGs reduce input size exponentially, it is possible that the problem of searching for Nash equilibria may become hard.

In the next section, we will show how to find the optimal Nash equilibrium for a certain class of optimality criteria in polynomial time. We will also exhibit optimality criteria for which the problem of finding optimal NE is $\mathcal{NP}$-hard.

Finding a Socially Optimal Nash Equilibrium

For a variety of optimality conditions, we can use dynamic programming to find an optimal NE in polynomial time.[2] We will first illustrate this technique by describing an algorithm to find the socially optimal NE.

Theorem 1. Opt-Nash *(Algorithm 1) calculates a socially optimal NE and terminates in* $O(n^6k^5)$ *time.*

Before proving this theorem, we will introduce some intuition and a pair of useful lemmas.

To find a socially optimal NE using dynamic programming, we will construct a Nash equilibrium from two "smaller" equilibria. For any subset of resources $T \subset R$ consider the *restricted* game $C|_T$ obtained by deleting resources in $R \setminus T$. Similarly, given a configuration ν, construct a restricted configuration $\nu|_T$ in $C|_T$ by ignoring players that do not choose resources in T. We exploit one key observation: if ν is a Nash equilibrium for C, then for any $T \subset R$, $\nu|_T$ is a NE for $C|_T$. The converse is not true: merging two NE configurations on different subsets of resources may not yield a NE, because a player may be able to switch from an expensive resource in its subset to an inexpensive resource in the other subset.

Thus, we need to enforce an additional constraint in merging two NEs. Let the *maximum exposed cost* $M(\nu, C)$ of a configuration be the cost of the most expensive resource that any player is using. Similarly, let the *minimum vacant cost* $V(\nu, C)$ be the cheapest cost a player could pay if he were to switch resources. We can merge two NE configurations on different subsets of resources only if the maximum exposed cost of each configuration is no more than the minimum vacant cost of the other.

[2] For monotonic costs, a faster greedy solution also works.

Algorithm 1 *Opt-Nash* (R, C, n)

// ***Dynamic Programming Table***: $A(h, B, M, V)$
// the cost of (B, M, V)-restricted-optimal NE
// on the set of resources $\{r_1, r_2, \ldots, r_h\}$.

// *Opt-Nash* ***Base Case:***

$$Z(i,B,M,V) = \begin{cases} B \cdot c(i,B) & \text{if } c(i,B) \le M \\ & \text{and } c(i,B+1) \ge V \\ \infty & \text{otherwise} \end{cases}$$

$A(1, B, M, V) = Z(1, B, M, V)$

for $h = 2$ to $|R|$ **do**

$A(h, B, M, V) =$

$$\min_{\substack{p \in \{0,1,\ldots,B\} \\ m \in \{C(i,j) | C(i,j) \le M\} \\ v \in \{C(i,j) | V \le C(i,j)\}}} \begin{array}{l} A(h-1, p, m, v) + \\ Z(h, B-p, \min\{M, v\}, \max\{V, m\}) \end{array}$$

end for

Return $A(|R|, n, \max\{C(i,j)\}, \min\{C(i,j)\})$.

Formally, assuming $C(\cdot, 0) = -\infty$ and $C(\cdot, n+1) = \infty$, define $M(\nu, C) = \max_r C(r, \nu_r)$ and $V(\nu, C) = \min_r C(r, \nu_r + 1)$.

Lemma 1. *Suppose T_1 and T_2 are disjoint subsets of R. Let $\nu|_{T_1}$ and $\nu|_{T_2}$ be two NE configurations for the restricted cost matrices $C|_{T_1}$ and $C|_{T_2}$. Furthermore, let $M(\nu|_{T_1}, C) \le V(\nu|_{T_2}, C)$ and $M(\nu|_{T_2}, C) \le V(\nu|_{T_1}, C)$. Then $\nu|_{T_1 \cup T_2} = \nu|_{T_1} \cup \nu|_{T_2}$ is also a NE for $C|_{T_1 \cup T_2}$.*

Proof. For contradiction, suppose $\nu|_{T_1 \cup T_2}$ is not a NE. Then there is a player p currently using resource i that is more expensive than another resource j. Since both $\nu|_{T_1}$ and $\nu|_{T_2}$ are NEs, assume wlog that $i \in T_1$ and $j \in T_2$.

By definition, $M(\nu|_{T_1}, C) \ge C(i, \nu|_{T_1}(i))$ and $V(\nu|_{T_2}, C) \le C(j, 1 + \nu|_{T_2}(j))$. Since p can move from i to j, $C(i, \nu|_{T_1}(i)) > C(j, 1 + \nu|_{T_2}(j))$. Thus, $M(\nu|_{T_1}, C) > V(\nu|_{T_2}, C)$, contradicting our assumption. that $M(\nu|_{T_1}, C) \le V(\nu|_{T_2}, C)$. □

Lemma 1 provides the basis for the merging step of our algorithm. We now introduce the notion of a *restricted-optimal* Nash equilibrium, and show that a socially optimal NE can be decomposed into smaller restricted-optimal NEs.

Definition 2. *A configuration ν for a cost matrix C is a* restricted-optimal *NE of parameters (B, M, V) if ν is socially optimal among all NEs ν' with B players where $M(\nu', C) \le M$ and $V(\nu', C) \ge V$.*

Lemma 2. *Let ν be a socially optimal NE for cost matrix C on the resources set R. Suppose $T \subset R$. Then $\nu|_T$ is a restricted-optimal NE of parameters* $\left(\sum_{i\in T} \nu_i, M(\nu|_T, C|_T), V(\nu|_T, C|_T)\right)$.

Proof. By contradiction, suppose $\nu|_T$ is not a restricted-optimal NE of the given parameters. Since $\nu|_T$ satisfies the parameters by definition, there exists a restricted-optimal NE ν' on T, which has a better social cost. Consider $\nu^* = \nu' \cup \nu|_{R\setminus T}$, which is a NE by Lemma 1. However, ν^* has a lower social cost than ν, contradicting our assumption of optimality. □

The algorithm *Opt-Nash* finds a socially optimal Nash configuration by adding one resource at a time to a restricted-optimal Nash equilibrium that only uses the first h resources. Using Lemmas 1 and 2, *Opt-Nash* combines restricted-optimal NEs on T_1 and T_2 for all possible combinations of (B, M, V) parameters to find the optimum. Algorithm 1 gives the pseudo code of *Opt-Nash* that returns the optimal value. With additional data structures, the actual assignment of players to resources can be obtained.

Proof of Theorem 1. Correctness of *Opt-Nash* is immediate from Lemma 2. The running time is determined by the number of entries in the table and the time it takes to fill a single entry. There are at most nk different values for M and V because the cost matrix only has at most nk entries. The value of B varies from 0 to n. We fill entries for $h = 2$ to k. Thus, there are a total $O(n^3k^3)$ table entries. For each entry, step 5 of *Opt-Nash* takes $O(n^3k^2)$ time, for the total of $O(n^6k^5)$. □

Note that the dynamic programming approach can be generalized to finding NEs that satisfy other interesting optimality criteria besides social welfare, so long as an appropriate version of Lemma 2 holds. For example, consider *minimax*-optimal Nash equilibrium, which minimizes the maximum cost that any player has to pay. To adapt *Opt-Nash* to find minimax-optimal NE, simply take the maximum of the two subproblems instead of their sum in step 5 of *Opt-Nash*.

Hard Optimality Criteria

In contrast to our previous result, there are optimality criteria for which finding the optimal Nash equilibrium is $\mathcal{NP}$-hard.

Definition 3. *A Nash configuration ν for a cost matrix C is a B-*bounded optimal *NE if for any Nash configuration ν' with $B \leq K(\nu', C)$, $B \leq K(\nu, C) \leq K(\nu', C)$.*

This is a natural criterion e.g. in a setting where some contractor hires coders to work on a number of independent projects simultaneously. If the workers are allowed to switch projects, it is natural to assume that they will settle in a Nash equilibrium. However, if the contractor has a budget of B, then only a B-bounded equilibrium will be acceptable.

Theorem 2. *Given n players, r resources, and the cost matrix C, it is $\mathcal{NP}$-hard to find a B-bounded optimal NE.*

Sketch. The theorem can be shown by a reduction from the PARTITION problem. □

Better- and Best-Response Dynamics

So far, we have concentrated on finding Nash equilibria using a centralized algorithm. While centrally computed optimal NEs carry normative power, in a multiagent system self-interested players would switch between resources on their own to lower their costs. This behavior is more appropriately captured by studying the dynamics of players making better or best responses. The fact that these natural dynamics converge to a NE in congestion games also gives the equilibrium concept a meaningful descriptive weight. In this section, we will quantify the rate at which Nash dynamics converges.

Note that there is no guarantee on the social cost of the NE to which Nash dynamics converges. There are examples where unless the agents start in a NE with low social cost, BR dynamics will always converge to a NE that is a factor of $\Theta(nk)$ more costly than the social optimum.

Two things were unspecified in the definition of Nash dynamics, both of which could possibly affect our analysis. First, Nash dynamics does not prescribe which player gets to move from a current state. Second, the initial strategy profile must be specified. Our theoretical results assume a worst-case player ordering and initial profile.

Upper Bounds

The first bound that we give shows that better response, and therefore best response, will converge in polynomial time.

Theorem 3. *For n players and k resources, better (hence, best) response converges in at most n^2k moves.*

The fact that Theorem 3 holds for better response is significant. It implies that even boundedly-rational agents that always take the first beneficial move will converge to a NE quickly.

To prove Theorem 3, observe that the dynamics of reaching a NE depends solely on preferences of the players between configurations rather than the actual costs.

Definition 4. *Given a cost matrix C, the* rank *of cost entry $C(i, j)$ is r if there are exactly $r - 1$* distinct *cost values (not entries) in C that are less than $C(i, j)$.*

Proof. Given an SCG cost matrix C, rank-preserving transformations of C do not affect any Nash equilibria or dynamics of reaching Nash equilibria. Therefore, replace each cost entry with its rank. The new potential function Φ can only have integral values between 1 and n^2k, since the highest rank is at most nk and there are n players. Any better-response move reduces Φ by at least 1, thus converging in n^2k moves. □

Our second bound limits the number of resources that are touched by any best-response dynamics in an SCG.

Theorem 4. *In a game with n players, no more than $2n$ resources are used during best-response dynamics.*

Sketch. Call player p a *pioneer* if p is the first player to use a resource for the first time. Under best response each player can be a pioneer at most once. □

Lower Bounds

We have examples that establish these lower bounds:

Theorem 5. *For n players and k resources,*

1. *There is a cost matrix C, an initial configuration ν, and a sequence of* better-response *moves $\mathcal{P}$ s.t. $|\mathcal{P}| = \Omega(n^{\frac{3}{2}}k)$.*

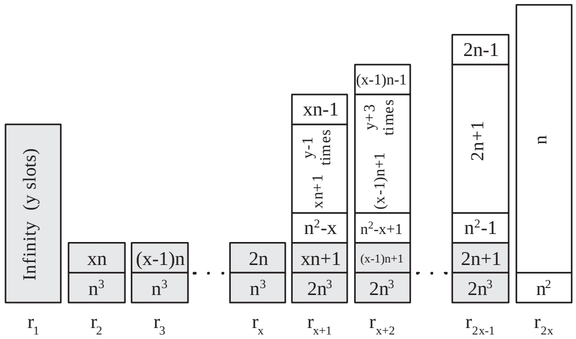

Figure 1: Best response takes $\Omega(\min(nk, n^2))$ moves.

2. *There is a cost matrix C and an initial configuration ν s.t. any* best *response sequence takes* $\Omega(\min(nk, n^2))$ *moves.*

Proof. Due to space constraints, we only illustrate case 2. Consider the cost matrix C in Figure 1. Each vertical column represents a resource. The cost of having j players on each resource is indicated by the values in the column, with the cost of having one player at the bottom of the column. Any unspecified costs have a value of infinity. The initial placement ν is represented by shaded boxes in the matrix.

Note, C has a special property: at each point only players using a particular resource have incentives to switch to a different resource. Initially, only players using r_1 can move to improve their payoffs. By tracing the best response dynamics, we see that all y players using r_1 will move to r_{x+1}. Then, the $xn-1$ entry of r_{x+1} becomes available. Now, a player using r_2 will move to improve his cost to $xn-1$. Note that r_{x+1} can only accommodate one of the two players from r_2, which means that the next player (now facing the cost of n^3) will move from r_2 to r_{x+2}. After this move, all players in r_{x+1} will move r_{x+2}. This pattern of all players moving from one resource to another will continue. Specifically, players using r_{x+i} will move to r_{x+i+1}. The total number of best response moves is $O(xy)$. By setting $y = \frac{1}{2}n$ and $x = \frac{1}{8}n$ we induce $\frac{3}{8}n^2$ moves. The $\min$ clause in the lower-bound is to reflect the fact that we actually need $O(n)$ resources to achieve the $O(n^2)$ lower-bound. $\square$

Simulation

In the previous section, we studied the worst-case behavior of BR dynamics in SCGs. To better understand the behavior of BR dynamics under more general circumstances, we conducted a number of simulation runs on random cost matrices.[3] In all of our experiments, we observed that BR dynamics converges rapidly to a NE. Due to space constraint, we will only highlight one of our experiments and speak briefly about the others.

In one experiment, we attempted to understand the effect of initial placement and choice of player-ordering heuristic on the rate of convergence. We generated a random cost matrix with 12 players and 6 resources and enumerated all 6188 possible initial configurations. We simulated BR dynamics for each initial configuration for three player-ordering heuristics: player with highest cost moving first, players moving in a round-robin fashion, and players moving in random order. The random heuristic was run 20 times with different seeds on each initial configuration. Figure 2 shows the CDF of the number of runs taking less than a given number of moves to converge to a NE. In all cases far fewer than $nk = 72$ moves, the theoretical lower bound, were required for convergence to a NE. Other heuristics and additional runs on other random cost matrices exhibited qualitatively similar results. In other experiments, we tested how well best-response dynamics scales with the size of the game. The results show that the number of moves required to converge to a NE grows extremely slowly as the number of players and resources increases.

[3] Random matrices provide the most unbiased view of BR behavior, even if they may not be typical of what happens in practice.

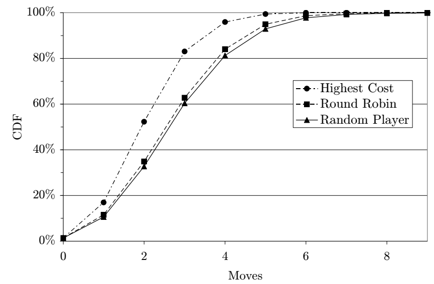

Figure 2: CDF of the number of moves, for a random 12-player 6-resource singleton game.

Independent-Resource Congestion Games

We now consider *independent-resource congestion games* (IRCGs), a natural extension of SCGs with much greater expressivity but similar theoretical properties. In IRCGs, each player is given a set $R_i \subset R$, and may choose any subset of R_i of limited size:

Definition 5. *An* independent-resource congestion game *is a CG* $\langle N, R, (S_i)_{i\in N}, (c_j)_{j\in M}, (l_i, u_i, R_i)_{i\in N}, \rangle$ *such that* $0 \le l_i \le u_i \le |R|$ *and* $S_i = \{X \subseteq R_i : l_i \le |X| \le u_i\}$.

IRCGs can model a much greater range of phenomena than SCGs, as players are no longer anonymous and may choose multiple resources. However, IRCGs are related to SCGs through the potential function: for an IRCG G, a strategy profile s with player i using k_i resources *induces* a configuration ν for an SCG $\tilde{G}$ that has $\sum_i k_i$ players and the same cost matrix, simply by treating each player as a set of k_i players each using one resource. After this mapping, $\Phi_G(s) = \Phi_{\tilde{G}}(\nu(s))$.

Given above, it is natural to hope for a tight link between Nash dynamics in the two classes. In fact, we will use a transformation to SCGs to show that best response in IRCGs will converge in polynomial time. However, we will also show that better response in IRCGs may take exponentially many moves to converge.

Theorem 6. *There exists a sequence of better responses in IRCGs that takes exponentially many moves to converge even if there is only one player in the game.*

Proof. Let there be k resources, with $C(i, 1) = 2^{i-1}$. Let $l_1 = 0$ and $u_1 = k$. Suppose that the only player initially uses all resources. The sequence of moves where each successive move reduces the cost of the player by 1 is a valid sequence of better responses. □

Theorem 7. *BR dynamics converges to a Nash equilibrium in $n^2k^2 + nk$ moves from any starting configuration.*

Proof. Let G be the IRCG. Define $\tilde{G}$ to be an SCG with the same cost matrix plus a resource r_0, with $C(r_0, \cdot) = 0$. Let the number of players in $\tilde{G}$ be $\sum_{i \in N} u_i$. Map each configuration ν in G to a configuration $\tilde{\nu}$ in $\tilde{G}$ as follows: $\tilde{\nu}_j = \nu_j$ for $j \neq 0$, and $\tilde{\nu}_0 = \sum_{i \in N} u_i - \sum_{j \in R} \nu_j$.

Suppose that some player i has a best response in G from S to S'. Let $S_0 = S, \ldots, S_m = S'$ be a sequence of resource choices for which each successive choice of resources differ in exactly one resource. This sequence can be constructed by setting $S_j = S_{j-1} - r'_j + r_j$, where $r_j \in S_{j-1} \setminus S'$ and $r'_j \in S' \setminus S_{j-1}$. In other words, choose a resource that does not belong to the final choice and replace with one that does. If $S_{j-1} \setminus S'$ or $S' \setminus S_{j-1}$ is empty, choose the appropriate resource to be r_0. Consider the configurations $\nu_0, ..., \nu_m$ in which player i plays S_j. We will show that each configuration ν_j has potential no higher than that of ν_{j-1}. Suppose that $\Phi_G(\nu_j) > \Phi_G(\nu_{j-1})$. Let ν' be the configuration where i plays $S' - r'_j + r_j$. Then $\Phi_G(\nu') < \Phi_G(\nu_m)$, contradicting that S' is a best response.

Now, $\tilde{\nu}_0, ..., \tilde{\nu}_m$ is a series of configurations in $\tilde{G}$ in which only one player moves and in which the potential is non-increasing. Furthermore, $\Phi_{\tilde{G}}(\tilde{\nu}_m) < \Phi_{\tilde{G}}(\tilde{\nu}_0)$. Thus, after a transformation from cost values to ranks, $\Phi_{\tilde{G}}(\tilde{\nu}_m) < \Phi_{\tilde{G}}(\tilde{\nu}_0)$. There are at most $nk + 1$ distinct ranks in $\tilde{G}$ and at most nk players. Thus, the potential is bounded by $n^2k^2 + nk$, and each best response in G decreases potential by at least 1. □

Note that a restricted form of better-response dynamics, which we term *local-improvement response dynamics*, where each agent makes local changes to its resource choice by replacing a resource used with a cheaper resource, will also converge to a NE in $n^2k^2 + nk$ moves.

Finally, we note that if for all agents i, $l_i = u_i = m$ for some constant m and $R_i = R$, *Opt-Nash* can be used to find an optimal NE; the complexity for general IRCGs remains an open question.

Conclusions and Open Questions

We introduced compactly-representable classes of singleton and independent-resource congestion games. We demonstrated that optimal Nash equilibria can be found in polynomial time in singleton games. We also gave force to the equilibrium concept by showing that better-response dynamics finds equilibria in polynomial time. Nevertheless, this work leaves a number of very important open questions.

First, we conjecture that the upper bound of n^2k on best response is too weak, and that an upper bound of $O(nk)$ is possible. However, proving this will require fundamentally new techniques: our analysis of Nash dynamics is based on potential functions, but potentials cannot easily differentiate between best and better responses, since they must decrease with every move. Thus, non-potential-based analysis of congestion games might be very useful.

We also feel that the area of compact forms for congestion games is barely explored, despite its considerable theoretical interest. One natural restriction would be to consider only contiguous resources (e.g. resources are temporal). If each agent must choose exactly m adjacent resources, we can find NE in $O(n^{m+1}k)$ time. Analysis of other compact, yet tractable classes will lead to more practical multiagent systems and lay sound theoretical ground for computational game theory.

Acknowledgements

Eugene Nudelman, Samuel Ieong, Robert McGrew, and Yoav Shoham are supported by NSF ITR-0205633. Robert McGrew is also supported by an NSF fellowship, and Samuel Ieong by a Stanford Graduate Fellowship. Qixiang Sun is supported by NSF CCR-0208683.

References

Anshelevich, E.; Dasgupta, A.; Kleinburg, J.; Tardos, E.; Wexler, T.; and Roughgarden, T. 2004. The price of stability for network design with fair cost allocation. In *FOCS04*, 295–304.

Bhat, N. A. R., and Leyton-Brown, K. 2004. Computing Nash equlibria of action-graph games. In *Proc. UAI-2004*.

Conitzer, V., and Sandholm, T. 2003. Complexity results about Nash equilibria. In *Proc. IJCAI-2003*, 765–771.

Even-Dar, E.; Kesselman, A.; and Mansour, Y. 2003. Convergence time to nash equilibria. In *ICALP2003*, 502–513.

Fabrikant, A.; Papadimitriou, C. H.; and Talwar, K. 2004. The complexity of pure Nash equilibria. In *Proc. STOC-2004*, 604–612.

Holzman, R., and Law-Yone, N. 1997. Strong equilibrium in congestion games. *Games and Economic Behavior* 21:85–101.

Kearns, M.; Littman, M.; and Singh, S. 2001. Graphical models for game theory. In *Proc. UAI-2001*.

Koller, D., and Milch, B. 2001. Multi-agent influence diagrams for representing and solving game. In *IJCAI-2001*.

Leyton-Brown, K., and Tennenholtz, M. 2003. Local-effect games. In *Proc. IJCAI-2003*.

Milchtaich, I. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13:111–124.

Osborne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory*. Cambridge, Massachusetts: The MIT Press.

Papadimitriou, C. 2001. Algorithms, games, and the internet. In *Proc. STOC-2001*, 749–753.

Porter, R.; Nudelman, E.; and Shoham, Y. 2004. Simple search methods for finding a Nash equilibrium. In *Proc. AAAI-2003*.

Rosenthal, R. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2:65–67.

Voorneveld, M.; Borm, P.; van Megan Freek; Tijs, S.; and Facchini, G. 1999. Congestion games and potentials reconsidered. *International Game Theory Review* 1:283–299.