

The Complexity of Welfare Maximization in Congestion Games

Carol A. Meyers*

Andreas S. Schulz**

June 2008; revised February 2010

Abstract

We investigate issues of complexity related to welfare maximization in congestion games. In particular, we provide a full classification of complexity results for the problem of finding a minimum cost solution to a congestion game, under the model of Rosenthal. We consider both network and general congestion games, and we examine several variants of the problem concerning the structure of the game and the properties of its associated cost functions. Many of these problem variants turn out to be NP-hard, and some are hard to approximate to within any finite factor, unless $P = NP$. We also identify several versions of the problem that are solvable in polynomial time.

1 Introduction

The study of *congestion games* was initiated by Rosenthal [24] as a simple class of games possessing pure-strategy Nash equilibria. The basic setup is as follows: we are given a finite number of players, each of which possesses a finite set of strategies. Each strategy consists of a subset of a master set of resources. The cost of employing a particular strategy is the sum of the costs of the resources associated with that strategy, where the cost of using a particular resource is solely a function of the number of players using that strategy. The cost of a resource is zero if it is not used.

One example of a congestion game occurs when the set of strategies is associated with paths in a network. In a *network congestion game*, each player i is associated with two nodes s_i and t_i , and the corresponding set of strategies consists of all (simple) $s_i - t_i$ paths. The arcs play the role of the resources, and the cost associated with each arc is a function of the number of players using that arc.

Rosenthal proposed two practical applications of congestion games, one concerning road networks and the other involving factory production. In the first application, a network of roads is given and each player travels from a certain origin to a certain destination. The cost of traveling on each road is an increasing function of the number of people using that road (hence the word ‘congestion’). In the second application, a number of firms are engaged in production, each of which has several production processes available that employ different resources. The cost of using a resource is a function of the number of firms that use the resource. Rosenthal showed that regardless of the cost structure on the set of resources, such games always possess a pure-strategy Nash equilibrium.

*Lawrence Livermore National Laboratory, L-153, 7000 East Avenue, Livermore, CA 94550

**Massachusetts Institute of Technology, E53-357, 77 Massachusetts Avenue, Cambridge, MA 02139

In this work we consider the issue of *welfare maximization* in congestion games, investigating in particular the complexity of congestion games from a *system optimal perspective*. This differs from the majority of the literature in congestion games, which typically address the existence and difficulty of finding Nash equilibria, and the various properties of such equilibria (see Section 2 for a literature review). Specifically, we address the complexity of finding an *overall minimum cost solution* to the congestion game problem, which is not required to be a Nash equilibrium.

One motivation for classifying the complexity of finding a minimum cost solution is that in some cases, we may be less interested in the performance of individual players than we are in the system optimum. This occurs in situations where a central planner is tasked with determining allocations to minimize global cost, such as in congestion games associated with load balancing or distributions of bandwidth. It can also occur in situations where the players are not selfish (i.e., if the players are all working together), but in which congestion effects are still experienced.

Other motivations for investigating welfare maximization include one suggested by the work of Anshelevich et al. [1], who showed how to obtain a provably good Nash equilibrium in certain cases starting from an optimal solution; for this method to be relevant in practice, we must know the complexity of finding such a solution. A final motivation is that in some problems, such as network congestion games with a single source and sink and nondecreasing costs, finding a Nash equilibrium may be done efficiently (see [9]) while computing the optimum is NP-hard (as we will show). In such cases, an algorithm for finding a Nash equilibrium may be used as an approximation algorithm for the problem of finding a minimum cost solution.

In what follows, we present our results on the complexity of finding system-optimal solutions to network and general congestion games. In Section 3, we introduce variants of the problems differing in *structure* and the type of associated *cost functions*. With regards to structure, we consider whether all players have the same set of strategies (symmetric) or not (asymmetric). In the network case, we also consider whether players have the same source or sink. For arc costs, we consider five different types (nondecreasing, convex nondecreasing, nonincreasing, concave nonincreasing, and nonmonotonic) that model different forms of congestion and economies of scale.

We fully categorize the complexity of the network congestion game problem and all of its variants under these parameters in Section 4. In most cases, we find the problem is NP-hard; however, in four cases (symmetric games with convex nondecreasing, nonincreasing, or concave nonincreasing arc costs, and single source games with concave nonincreasing arc costs) the problem is solvable in polynomial time. We examine the complexity of the general congestion game problem in Section 5. In several cases, our results follow directly from the network case, but in others (for instance, convex nondecreasing costs) we are able to derive stronger results. Overall, we find that in almost all cases, the problems are NP-hard and difficult to approximate to any finite factor, unless $P = NP$.

2 Related Work

The topic of welfare maximization in congestion games was recently and independently examined in two different studies. In both of these works, the (notably different) model of congestion games due to Milchtaich [19, 20, 21] was used, where players possess player-specific cost functions and travel over parallel links. In this model, the effects of congestion are associated with the players rather than the resources; each player experiences a different *player-specific* amount of congestion according to the number of other players sharing the resources it uses.

In the first study, Chakrabarty et al. [5] showed that finding a minimum cost solution in congestion games under the Milchtaich model is NP-hard and no finite-factor approximation algorithm exists (unless $P = NP$). They showed that in certain special cases, such as where all of the strategies cost the same and the matrix of player costs is anti-Monge, the system optimum may be computed in polynomial time. In the second study, Blumrosen and Dobzinski [4] showed that there exists an 18-approximation algorithm for welfare maximization in congestion games under the Milchtaich model, regardless of the cost structure. In the case where the costs are nondecreasing in the number of players sharing the resource, the performance ratio can be improved to $\frac{e}{e-1} \approx 1.582$. They also demonstrated a connection between such congestion games and combinatorial auctions.

Our work differs from that of Chakrabarty et al. and Blumrosen and Dobzinski in that we consider the Rosenthal [24] model of congestion games, rather than the Milchtaich [19] model. We also consider a greater number of structural and combinatorial aspects of the problem and a variety of different resource cost functions, as we describe in Section 3.

The study of welfare maximization is complementary to the larger body of literature in congestion games, which has primarily focused on the existence and complexity of finding *Nash equilibria* in such games, and the price of anarchy of such equilibria. In particular, much work [6, 15, 18, 25] has considered the properties of equilibria in games consisting of n users traveling over m parallel arcs, the study of which was initiated by Koutsoupias and Papadimitriou [17]. An excellent survey of results in this area is given by Kontogiannis and Spirakis [16]. Returning to general games, Fotakis, Kontogiannis, and Spirakis [11, 12] showed that equilibria may not always exist in weighted congestion games, where each player may control different amounts of demand; Beier et al. [3] addressed the problem of computing equilibria in games with imperfect information.

The study of complexity issues in congestion games was initiated by Fabrikant, Papadimitriou, and Talwar [9]. They showed that a pure Nash equilibrium can be computed in polynomial time in network congestion games with nondecreasing arc costs where all players share a source and sink; however, in general the problem is PLS-complete. Following this, Jeong et al. [14] showed that in congestion games with parallel links and general costs, a best Nash equilibrium can be computed in polynomial time. In addition, Papadimitriou and Roughgarden [22, 23] demonstrated that a generalization known as a correlated equilibrium can be calculated in polynomial time in compactly encoded games. In terms of hardness, Dunkel and Schulz [7] showed that it is strongly NP-hard to determine whether a weighted congestion game possesses a Nash equilibrium.

3 Preliminaries

In a *general congestion game*, we are given a set of *resources* $A = \{a_1, \dots, a_m\}$ and a set of *players* $P = \{1, \dots, n\}$. Each player i possesses a set of allowable *strategies*, where each such strategy $s_i \subseteq A$ consists of a subset of the resources. Each player wishes to select and play exactly one strategy. A *solution* $s = (s_1, \dots, s_n)$ consists of the chosen strategies for each player.

The *cost of a resource* $a \in A$ is given by a nonnegative function $c_a(x)$ that computes the per-unit cost of x players using a . The cost function may be arbitrary in general, but it is restricted to being solely a function of the number of players using the resource. The *cost of a strategy* s_i is the sum of the costs of the resources associated with that strategy. The *cost of a solution* s is equal to

$$\sum_{a \in A} x_a c_a(x_a),$$

where $x_a = |\{i : a \in s_i\}|$ is the total number of players using resource a in the solution, and s_i is the strategy chosen by player i in s .

A *network congestion game* is a special case of a general congestion game in which resources are associated with *arcs*, strategies are associated with *simple paths*, and players are associated with *units of demand* in a network. This is a special type of minimum cost integer multicommodity flow problem where the cost per unit flow on each arc differs based on how much flow is traversing the arc. On the other hand, there are no arc capacities.

More formally, in a network congestion game we are given a directed graph $G = (N, A)$ and a set of players $P = \{1, \dots, n\}$. Each player i is associated with a pair of nodes $s_i \in N$ and $t_i \in N$, with the understanding that player i wishes to send one unsplittable unit of flow from node s_i to node t_i . If an arc $a = (u, v)$ is labeled as $c_a(1)/c_a(2)/c_a(3)/\dots/c_a(n)$, then the cost of sending 1 unit of flow along the arc is $c_a(1)$, the cost of sending 2 units of flow is $c_a(2)$ per unit (for a total cost of $2 \cdot c_a(2)$), and the cost of sending k units of flow is $c_a(k)$ per unit (for a total cost of $k \cdot c_a(k)$). Each arc has n different labels, since n is the greatest number of players that can traverse an arc. The goal is to route each player on a single path from its source to its sink in a minimum cost manner.

We consider several variants of these problems. In terms of structure, we consider two basic alternatives: *symmetric* problems, in which all players share the same set of strategies, and *asymmetric* problems, where players may have different strategies. In the network problem, the symmetric case corresponds to all players having the same source and sink, and the asymmetric case corresponds to different sources and sinks. In addition, in the network problem we also consider the *single source* case, in which all players share a single source (but may have different sinks).

With regards to cost functions, we consider five different classes of cost structures. We say that the arc costs are *nondecreasing* if $c_a(1) \leq c_a(2) \leq \dots \leq c_a(n)$ for all $a \in A$, and *nonincreasing* if $c_a(1) \geq c_a(2) \geq \dots \geq c_a(n)$ for all $a \in A$. Nondecreasing cost functions model the negative effects of congestion on resource availability, while nonincreasing cost functions reflect economies of scale. We say that a cost structure is *convex nondecreasing* if the differences between consecutive aggregate arc costs are nondecreasing; in other words, $ic_a(i) - (i-1)c_a(i-1) \leq (i+1)c_a(i+1) - ic_a(i)$ for all $i = 1, \dots, n-1$. Similarly, we say a structure is *concave nonincreasing* if $ic_a(i) - (i-1)c_a(i-1) \geq (i+1)c_a(i+1) - ic_a(i)$ for all $i = 1, \dots, n-1$.¹ Note that convex nondecreasing (concave nonincreasing) arc costs are automatically nondecreasing (nonincreasing) as well. If an arc cost function fits into none of these categories, we say that it is *nonmonotonic*.

This gives us fifteen different problems in the network case (three structural variants and five cost variants), and ten different problems in the general case (two structural variants and five cost variants). As it turns out, many of the complexity results we prove apply to multiple problems, with minor changes.

4 Network Complexity Results

Our complexity results for network congestion games are as illustrated in Table 1.

We cover the hardness results first: we begin by presenting the single source hardness results, and we show how slight modifications can be made to derive the symmetric hardness results. We then give the asymmetric hardness results, followed by the polynomial time algorithms.

¹Put differently, a per-unit arc cost function $c_a(i)$ is convex nondecreasing (concave nonincreasing) if and only if the total cost function $i \cdot c_a(i)$ is supermodular (submodular).

type→ costs ↓	symmetric	single source	asymmetric
nondecreasing	NP-hard [4.1]	NP-hard [4.1]	Inapprox.
convex nondecreasing	P	P [4.6]	Inapprox. [4.4]
nonincreasing	P [4.5]	Inapprox. [4.2]	Inapprox.
concave nonincreasing	P	NP-hard [4.3]	NP-hard
nonmonotonic	Inapprox. [4.2]	Inapprox.	Inapprox.

Table 1: Complexity results for network congestion games.

(The numbers in brackets indicate in which theorem (or discussion thereafter) the results are proved. Note that the results for all unlabeled entries follow directly from other entries on the table. By ‘Inapprox.’ we mean that the problem is NP-hard, and the reduction actually shows that it is NP-hard to decide whether the optimal cost is zero or not. Thus, no finite-factor approximation algorithm can exist for the respective problem, unless $P=NP$.)

Our first theorem concerns single source congestion games with nondecreasing costs. (Note that the hardness of this problem does not follow directly from the hardness of the general single-source unsplittable flow problem (see [2]), as the latter problem translates to *weighted* congestion games.)

Theorem 4.1 *The single source network congestion game problem with nondecreasing costs is strongly NP-hard.*

Proof: We reduce from 3-PARTITION, which is strongly NP-complete [13]. This problem is:

Instance: A set $R = \{r_1, \dots, r_{3q}\}$ of $3q$ elements, a bound $B \in \mathbb{Z}^+$, and a size $z(r_i) \in \mathbb{Z}^+$ for each $r_i \in R$, such that $\frac{B}{4} < z(r_i) < \frac{B}{2}$ and $\sum_{r_i \in R} z(r_i) = qB$.

Question: Can R be partitioned into q disjoint sets D_1, \dots, D_q such that, for all $1 \leq j \leq q$, we have $\sum_{r_i \in D_j} z(r_i) = B$?

Suppose we are given an instance of the 3-PARTITION problem. Build the following congestion game (see Figure 1):

1. Create 1 source node s , with a supply of $qB + 3q^2$.
Create $3q$ transshipment nodes s_i , where s_i corresponds to element r_i in R .
Create q sink nodes D_j , each with demand B .
Create $3q^2$ sink nodes a_{ij} , $1 \leq i \leq 3q$ and $1 \leq j \leq q$, each with demand 1.
2. Add arcs (s, s_i) of cost $0/\dots/0/M/\dots/M$, where the last ‘0’ is in place $z(r_i)+q$ and M is a large number.
Add arcs (s_i, a_{ij}) of cost $0/1/\dots/1$, for all i, j satisfying $1 \leq i \leq 3q$ and $1 \leq j \leq q$.
Add arcs (a_{ij}, D_j) of cost $0/0/\dots/0$, for all i, j satisfying $1 \leq i \leq 3q$ and $1 \leq j \leq q$.

In terms of the game structure, this corresponds to $qB + 3q^2$ players having origin s , B players having destination D_j , and one player having destination a_{ij} , for all i and j .

We claim that if the 3-Partition answer is ‘yes,’ then the congestion game cost is equal to $qB + 3q$; if the answer is ‘no,’ the congestion game cost is greater than or equal to $qB + 3q + 1$. To see the first implication, suppose the sets in our 3-partition are D_1, \dots, D_q . We construct a solution as follows. First, send 1 unit of flow along each of the paths $s - s_i - a_{ij}$, for all $i \in \{1, \dots, 3q\}$ and $j \in \{1, \dots, q\}$. Next, for all $r_i \in R$ such that $r_i \in D_j$, send $z(r_i)$ units of flow along the path $s - s_i - a_{ij} - D_j$. This solution will be feasible, since D is a 3-partition and thus the inflow at each

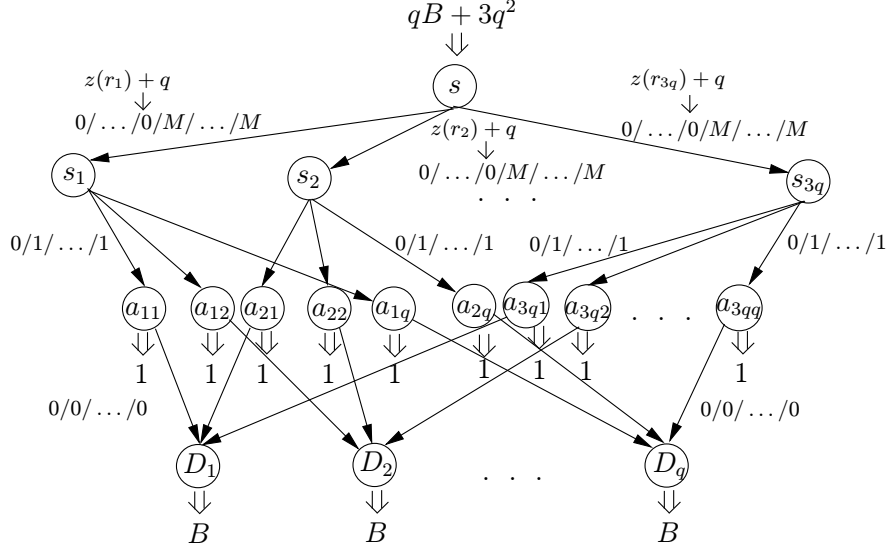


Figure 1: Constructed instance of the congestion game problem with nondecreasing arc costs.

node D_i will be equal to B . By inspection, the cost contributed by arc (s_i, a_{ij}) is equal to $z(r_i) + 1$ if $r_i \in D_j$, and 0 otherwise. Hence the total cost of the solution is $qB + 3q$.

To see the second implication, suppose there is a solution to the congestion game problem of cost at most $qB + 3q$. Because all arcs exiting from the nodes s_i have the cost structure $0/1/\dots/1$, it follows that no node s_i can have more than one unit of flow exiting along two arcs (s_i, a_{ij}) and $(s_i, a_{ij'})$, for $j \neq j'$. (Otherwise, the total cost would be strictly greater than $qB + 3q$, as there are $qB + 3q^2$ units of flow that enter these nodes.) This implies in particular that s_i must send $z(r_i) + 1$ units of flow along exactly one of the arcs (s_i, a_{ij}) , and consequently $z(r_i)$ units of flow travel from s_i to that corresponding node D_j . Consider the partition where each element r_i is mapped to the set D_j that its corresponding node s_j sends flow to in the solution. Each of the sets D_j will have size B , since the demand of each node D_j is B , and no element r_i will be mapped to more than one of the sets D_j . Hence this is a 3-Partition of the elements in R . \square

We can easily extend this result to the symmetric version of the problem: add a new node t and new arcs (a_{ij}, t) and (D_j, t) for all i, j . Set the cost of the new (a_{ij}, t) arcs to $0/1/\dots/1$ and the cost of the (D_j, t) arcs to $0/0/\dots/0/M/\dots/M$, where the last '0' occurs in the B -th position. Adjust the demand so that the nodes a_{ij} and D_j have demand 0 and the node t has a demand of $qB + 3q^2$. The same conclusions will hold.

A similar construction gives an even stronger result for single source games with nonincreasing arc costs. Recall that an α -approximation algorithm is a polynomial-time algorithm that produces a feasible solution of cost within a factor of α of the optimum.

Theorem 4.2 *The single source network congestion game problem with nonincreasing arc costs is strongly NP-hard, and it does not have a finite-factor approximation algorithm, unless $P=NP$.*

Proof: Again we reduce from the 3-PARTITION problem. We use a simplified version of the construction in Theorem 4.1. We build a graph as follows (see Figure 2):

1. Create 1 source node s , with a supply of $qB + 3q$.




Figure 2: Constructed instance of the congestion game problem with nonincreasing arc costs.

2. Create $3q$ sink nodes s_i , each with a demand of 1.
Create q sink nodes D_j , each with a demand of B .
3. Add arcs (s, s_i) of cost $M/M/\dots/M/0/\dots/0$, for all i , where the first '0' occurs in the $(z(r_i) + 1)$ -st place, and M is a large number.
Add arcs (s_i, D_j) of cost $M/M/\dots/M/0/\dots/0$, for all i, j , where the first '0' occurs in the $z(r_i)$ -th place.

In terms of the game framework, this corresponds to $qB + 3q$ players having origin s , 1 player having destination s_i , and B players having destination D_j , for all i and j .

If the answer to the 3-PARTITION problem is 'yes,' we can obtain a routing of cost 0 by sending $z(r_i) + 1$ units of flow from the source s to node s_i , and then routing $z(r_i)$ of those units from s_i to the node D_j corresponding to the set r_i is mapped to in the partition. Conversely, if the optimal solution to the constructed instance of the congestion game problem has cost 0, exactly $z(r_i) + 1$ units of flow are sent on each arc (s, s_i) , and exactly $z(r_i)$ units of flow are sent along one arc (s_i, D_j) . Thus, we can obtain a 3-partition of the elements r_i by placing each element r_i into the set D_j that s_i sends flow to in the optimal solution. There will be only one such set, because of the cost structure, and each set D_j will have size B , due to the way the demands are defined.

This implies that the single source network congestion game problem with nonincreasing arc costs is strongly NP-hard. In fact, the reduction shows that it is even hard to distinguish between instances of cost 0 and instances of cost greater than 0. This implies that we cannot approximate this problem to within any finite factor, unless $P=NP$. \square

The result can be extended to the symmetric problem with nonmonotonic arc costs: add a super-sink t as in the discussion following Theorem 4.1, and arcs (s_i, t) and (D_j, t) for all i, j . Set the costs of the arcs (s_i, t) to $0/M/\dots/M$, and the costs of the arcs (D_j, t) to $0/\dots/0/M/\dots/M$, where the last '0' is in the B th position. The same conclusions follow.

The same argument does not apply to concave nonincreasing arc costs, but a simple reduction gives that this problem is NP-hard as well.

Theorem 4.3 *The single source network congestion game problem with concave nonincreasing arc costs is strongly NP-hard.*

Proof: We reduce from the DIRECTED STEINER TREE problem, which is strongly NP-complete [13]. This problem is:

Instance: A directed graph $G = (N, A)$ with arc costs $c_a \in \mathbb{Z}^+$ for all $a \in A$, a root node s , a set of terminals $\{t_1, t_2, \dots, t_n\} \subseteq V$, and a bound $B \in \mathbb{Z}^+$.

Question: Does there exist a directed tree T rooted at node s , such that T contains an $s - t_i$ path for all $i = 1, \dots, n$, and the sum of the combined arc costs in T is at most B ?

Suppose we are given $G = (N, A)$, the node s , and terminals t_1, \dots, t_n . We define an instance of the congestion game on G as follows. First, we assign one player to travel from node s to node t_i , for all $i \in \{1, \dots, n\}$. Second, we set the cost of an arc $a \in A$ equal to the concave nonincreasing function $c_a / \frac{c_a}{2} / \frac{c_a}{3} / \dots / \frac{c_a}{n}$.

We claim that there is a solution to this single source congestion game problem of cost at most B if and only if there is a solution to the DIRECTED STEINER TREE problem of cost at most B . To see the first direction, suppose there exists a directed Steiner tree T of cost at most B . Since T is a Steiner tree, it must contain a path from s to t_i for all i . Consider a solution to the congestion game problem where we route all players from s to t_i using only arcs contained in T . This solution will be feasible since it contains a feasible path for every player, and its total cost will be at most B , by the way the costs in the congestion game are defined.

To see the other direction, suppose there exists a solution to the congestion game problem of cost at most B . The collection of arcs used in this solution must contain a path from s to t_i for all i , since the congestion game solution is feasible. Hence it must also contain a directed Steiner tree T . The cost of this tree will be at most B , by the way the congestion game costs are defined. Hence the congestion game problem is NP-hard. \square

Another relatively simple reduction provides a strong hardness result for the asymmetric congestion game problem with convex nondecreasing costs.

Theorem 4.4 *The asymmetric network congestion game problem with convex nondecreasing arc costs is strongly NP-hard, and it does not have a finite-factor approximation algorithm, unless $P=NP$.*

Proof: We reduce from the ARC-DISJOINT PATHS problem, which is strongly NP-complete [13]. This problem is:

Instance: A directed graph $G = (N, A)$ and a set of node pairs $(s_1, t_1), \dots, (s_n, t_n)$.

Question: Does there exist a collection of arc-disjoint paths P_1, \dots, P_n , where P_i is an $s_i - t_i$ path?

Suppose we are given $G = (N, A)$ and $(s_1, t_1), \dots, (s_n, t_n)$. We transform this into an instance of the asymmetric network congestion game problem with convex nonincreasing costs as follows. First, we assign one player to travel from node s_i to node t_i , for all i . Second, for every arc $(i, j) \in A$, we introduce the cost structure $0/1/2/\dots/n-1$.

If there exist arc-disjoint paths, then any routing using these paths will have cost 0, since each arc will be taken at most once. Conversely, if there do not exist arc-disjoint paths, in any routing some arc will have to be taken twice, for a cost of at least 1. Hence the asymmetric network congestion game problem with convex nondecreasing arc costs is NP-hard; moreover, it is already hard to distinguish between instances of cost 0 and those of cost greater than 0. \square

We have now covered all of the hardness results. We next address variants of the problem that are solvable in polynomial time.

Theorem 4.5 *The symmetric network congestion game problem with nonincreasing arc costs is solvable in polynomial time.*

Proof: Suppose we are given an instance of this problem, consisting of $G = (N, A)$, designated nodes s and t , nonnegative arc costs, and a collection of players $\{1, \dots, n\}$. We claim that in such an instance, there exists an optimal solution where all players follow the same path from their origin to their destination.

To see this, suppose in a solution at least two players follow different paths. Let c_i denote the cost of the path followed by player i . Further suppose that among all the players, player k is following a path of minimal cost c_k . Now, consider rerouting all of the other players onto the path followed by player k . Since the arc costs are nonincreasing, the cost of this path will change to $c'_k \leq c_k$. The total cost of the solution will change to $nc'_k \leq nc_k \leq \sum_i c_i$. Hence there is some optimal solution where all players follow the same path.

In a solution where all players follow the same path, the cost of each arc a in the solution is equal to the cost $c_a(n)$ of routing n players across the arc. This suggests a simple algorithm for solving the problem: first, fix the cost of each arc $a \in A$ to $c_a(n)$; next, find the shortest $s - t$ path in G with respect to the new arc costs, and route all n players along this path. This gives a minimum cost solution to the problem where all players follow the same path, so it is optimal. \square

We have one final complexity result, which relates to convex nondecreasing arc costs.

Theorem 4.6 *The single source network congestion game problem with convex nondecreasing arc costs is solvable in polynomial time.*

Proof: This result was independently proved by Chakrabarty et al. [5], though their proof was only stated for the symmetric case and linear costs. For completeness, we review the result here, noting that it also extends to the single source case.

Suppose we have an instance of the single source problem. We give a reduction to the minimum cost flow problem. We create a new graph G' on the same node set N , where the arcs are defined as follows. For every arc $a \in A$ with cost structure $c_a(1)/c_a(2)/\dots/c_a(n)$, we introduce n parallel arcs a_1, a_2, \dots, a_n in G' with the same head and tail nodes as a , where the cost of arc a_k is equal to $kc_a(k) - (k-1)c_a(k-1)$ and the capacity of each arc is 1.

We claim a minimum cost flow on G' gives a minimum cost flow on G , by setting the flow on $a \in A$ equal to the sum of the flows on the corresponding arcs in G' . To see this, first observe that there exists an integral minimum cost flow on G' , since standard network flow problems with integer capacities always admit an integral optimal solution. Moreover, because the costs are convex nondecreasing, there exists such an optimal solution so that any flow traveling across the parallel arcs a_1, \dots, a_n in G' will fill in order of increasing index (from 1 to n). This implies that if there are k units traveling across a set of parallel arcs, the corresponding cost will be

$$c_a(1) + (2c_a(2) - c_a(1)) + \dots + (kc_a(k) - (k-1)c_a(k-1)) = kc_a(k).$$

Thus the cost structure in G' mimics that of G , and it follows that a minimum cost flow in G' gives an integral minimum cost flow in G . \square

Fixed Number of Players We now comment on the complexity of the aforementioned problems with a fixed number of players. Here, we are given a set of players $\{1, 2, \dots, n\}$, where n is a fixed constant. (Thus the running time of a polynomial-time algorithm may depend exponentially on n .)

In general congestion games, where the set of strategies is given explicitly, the congestion game problem with a fixed number of players can be solved in polynomial time (every player tries every strategy). In the case of network congestion games, however, there may be an exponential number of strategies: here the strategy space is compactly encoded, and the number of strategies can be exponential in the number of nodes and arcs in the network. Thus for this problem, the case with a fixed number of players is a nontrivial variant.

Several of our earlier results can be extended to apply to network congestion games with a fixed number of players. In particular, Theorem 4.4 holds for a fixed number of players, since the ARC-DISJOINT PATHS problem is NP-complete even for only two terminal pairs [10]. Similarly, Theorems 4.5 and 4.6 apply, since anything that can be solved in polynomial time with an arbitrary number of players can be solved in polynomial time with a fixed number of players.

We also observe that the single source network congestion game problem with concave nonincreasing arc costs can be solved in polynomial time for a fixed number of players. This is since we can model the problem as a minimum cost integer network flow problem with a concave cost function, by taking a piecewise linearization of the arc costs. (In other words, we create a continuous cost function for the problem by fitting a straight line between each two consecutive specified arc costs.) Such problems can be solved in polynomial time for fixed demand using the Send-and-Split method proposed by Erickson, Monma, and Veinott [8].

5 General Complexity Results

Our complexity results for general congestion games are given in Table 2. We note that since the set of strategies in general congestion games is given *explicitly* rather than *implicitly*, this problem is different from the network case, and the same results do not immediately apply. In fact, the input size of an instance of a general congestion game is determined by the number of resources, the number of players, the number of strategies, and the maximum encoding length over all costs. In contrast, the input size of an instance of a network congestion game only depends on the number of players, the size of the network, and the maximum encoding length of an arc cost. As in the previous section, we first present the hardness results and then a polynomial time algorithm.

type→ costs ↓	symmetric	asymmetric
nondecreasing	Inapprox.	Inapprox.
convex nondecreasing	Inapprox. [5.1]	Inapprox.
nonincreasing	P [5.4]	Inapprox. [5.2]
concave nonincreasing	P	NP-hard [5.3]
nonmonotonic	Inapprox.	Inapprox.

Table 2: Complexity results for general congestion games

Theorem 5.1 *The symmetric general congestion game problem with convex nondecreasing arc costs is strongly NP-hard, and no finite-factor approximation algorithm exists unless $P=NP$.*

Proof: We reduce from the 3-DIMENSIONAL MATCHING problem, which is strongly NP-complete [13]. This problem is:

Instance: A set $S \subseteq X \times Y \times Z$, where X , Y , and Z are disjoint sets of cardinality q each.

Question: Does S contain a subset $S' \subseteq S$ such that $|S'| = q$ and no two elements of S' agree in any coordinate?

Suppose we are given X , Y , Z , and S . We define an instance of the general congestion game problem as follows: let the members of X , Y , and Z correspond to the resources, and let each $s \in S$ correspond to a potential strategy. (Thus, each strategy contains three resources: one from X , one from Y , and one from Z .) Define q players, each of which possesses the same set of strategies S . Set the cost of each resource to $0/1/2/\dots/(q-1)$.

We claim that if the answer to the 3-DIMENSIONAL MATCHING problem is ‘yes,’ then the optimal cost of the congestion game problem is 0; if the answer is ‘no,’ then the cost is at least 1. To see this observe that a solution to the congestion game problem has cost 0 if and only if the strategies chosen by players in that solution constitute a matching. Moreover, if there is no matching, some resource will have to be chosen more than once, for a cost of at least 1. \square

In the case of nonincreasing arc costs, our proof from the previous section carries over.

Theorem 5.2 *The asymmetric general congestion game problem with nonincreasing arc costs is strongly NP-hard and cannot be approximated to within any finite factor, unless $P=NP$.*

Proof: This follows directly from Theorem 4.2. Note that in our construction, there are a total of $3q^2$ possible strategies, which is polynomial in the input size. \square

For concave nonincreasing costs, we give a somewhat different argument.

Theorem 5.3 *The asymmetric general congestion game problem with concave nonincreasing arc costs is strongly NP-hard.*

Proof: We reduce from MINIMUM COVER, which is strongly NP-complete [13]. This problem is:

Instance: A finite set X , a collection S of subsets of X , and an integer $K \leq |S|$.

Question: Does S contain a subset $S' \subseteq S$ with $|S'| \leq K$, such that every element of X belongs to at least one member of S' ?

Suppose we are given X , S , and K . We construct an instance of the asymmetric general congestion game problem with concave nonincreasing arc costs as follows. Let the sets in S correspond to both the resources and the strategies, so that each strategy consists of one resource. Set the cost of each resource to $1/\frac{1}{2}/\frac{1}{3}/\dots/\frac{1}{n}$, where $n = |X|$. Define n players, each corresponding to an element of X , and set the possible strategies associated with player $x \in X$ to be those sets $S_x \subseteq S$ containing element x .

We claim that the answer to the MINIMUM COVER problem is ‘yes’ if and only if the optimal cost of the congestion game is less than or equal to K . To see this, first note that each resource $s \in S$ costs the same regardless of how many players are using it. Thus the optimal solution to the congestion game corresponds to the smallest collection of sets that cover all the elements in X . It follows that if the optimal cost is less than or equal to K , then the corresponding instance is a ‘yes’ instance of the problem. Conversely, if there is a minimum cover of size less than or equal to

K , we can obtain a solution to the congestion game of cost less than or equal to K by selecting for each element a strategy that contains it in the minimum cover. \square

Finally, we have one polynomial-time algorithm.

Theorem 5.4 *The symmetric general congestion game problem with nonincreasing arc costs is solvable in polynomial time.*

Proof: By a similar argument to that in Theorem 4.5, we see that in any such problem it is optimal for all players to choose the same strategy. Hence we need only determine the cheapest strategy, where the cost of each resource $a \in A$ is set to $c_a(n)$. This can be done in polynomial time, because the number of strategies is part of the input. \square

6 Concluding Remarks

We have provided the first extensive study of the complexity of finding welfare maximizing (or equivalently, minimum cost) solutions to congestion games, under the model of Rosenthal [24]. This complements the previous work on welfare maximization by Chakrabarty et al. [5] and Blumrosen and Dobzinski [4], who considered the Milchtaich [19] model of congestion games. For the most part, the problems we studied are NP-hard, but we have identified several variants that are solvable in polynomial time. We examined a variety of different structural aspects and several different types of cost functions.

We note that for those problems in which we showed a multiplicative approximation algorithm does not exist, allowing for an additive term does not help. This can be seen by multiplying all non-zero costs on the arcs by an arbitrarily large factor of M .

Finally, we have not yet addressed the approximability of those NP-hard problems contained in this paper for which inapproximability results were not obtained. This is an intriguing area for further study, as it may provide new insights into the structure and properties of this class of problems.

References

- [1] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the Forty-Fifth Symposium on Foundations of Computer Science*, pages 295–304, Rome, Italy, 2004. IEEE Computer Society.
- [2] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and their Approximability Properties*. Springer, Berlin, Germany, 1999.
- [3] R. Beier, A. Czumaj, P. Krysta, and B. Vöcking. Computing equilibria for a service provider game with (im)perfect information. *ACM Transactions on Algorithms*, 2:679–706, 2006.
- [4] L. Blumrosen and S. Dobzinski. Congestion games with malicious players. In *Proceedings of the Seventh ACM Conference on Electronic Commerce*, pages 52–61, Ann Arbor, MI, 2006. ACM Press.
- [5] D. Chakrabarty, A. Mehta, V. Nagarajan, and V. Vazirani. Fairness and optimality in congestion games. In *Proceedings of the Sixth ACM Conference on Electronic Commerce*, pages 52–57, Vancouver, BC, 2005. ACM Press.

- [6] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms*, 3, 2007.
- [7] J. Dunkel and A.S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research*, 33:851–868, 2008.
- [8] R. Erickson, C. Monma, and A. Veinott Jr. Send-and-split method for minimum-concave-cost network flows. *Mathematics of Operations Research*, 12:634–664, 1987.
- [9] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pages 604–612, Chicago, IL, 2004. ACM Press.
- [10] S. Fortune, J. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10:111–121, 1980.
- [11] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348:226–239, 2005.
- [12] D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in network congestion games: Pure equilibria and anarchy cost. In *Proceedings of the Third Workshop on Approximation and Online Algorithms*, number 3879 in Lecture Notes in Computer Science, pages 161–175, Mallorca, Spain, 2005. Springer.
- [13] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, New York, NY, 1979.
- [14] S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: a simple class of congestion games. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, pages 489–494, Pittsburgh, PA, 2005. AAAI Press.
- [15] R. Johari and J. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29:407–435, 2004.
- [16] S. Kontogiannis and P. Spirakis. Atomic selfish routing in networks: a survey. In *Proceedings of the First International Workshop on Internet and Network Economics*, number 3828 in Lecture Notes in Computer Science, pages 989–1002, Hong Kong, China, 2005. Springer.
- [17] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pages 404–413, Trier, Germany, 1999. Springer.
- [18] M. Mavronicolas and P. Spirakis. The price of selfish routing. *Algorithmica*, 48:91–126, 2007.
- [19] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.
- [20] I. Milchtaich. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research*, 25:349–364, 2000.
- [21] I. Milchtaich. The equilibrium existence problem in finite network congestion games. In *Proceedings of the Second International Workshop on Internet and Network Economics*, number 4286 in Lecture Notes in Computer Science, pages 87–98, Patras, Greece, 2006. Springer.
- [22] C. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, pages 49–56, Baltimore, MD, 2005. ACM Press.
- [23] C. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 82–91, Vancouver, BC, 2005. Society for Industrial and Applied Mathematics.

- [24] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [25] S. Suri, C. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47:79–96, 2007.