

The Price of Stability for Network Design with Fair Cost Allocation

Elliot Anshelevich* Anirban Dasgupta† Jon Kleinberg ‡ Éva Tardos§
Tom Wexler¶ Tim Roughgarden||

Abstract

Network design is a fundamental problem for which it is important to understand the effects of strategic behavior. Given a collection of self-interested agents who want to form a network connecting certain endpoints, the set of stable solutions — the Nash equilibria — may look quite different from the centrally enforced optimum. We study the quality of the best Nash equilibrium, and refer to the ratio of its cost to the optimum network cost as the price of stability. The best Nash equilibrium solution has a natural meaning of stability in this context — it is the optimal solution that can be proposed from which no user will “defect”.

We consider the price of stability for network design with respect to one of the most widely-studied protocols for network cost allocation, in which the cost of each edge is divided equally between users whose connections make use of it; this fair-division

scheme can be derived from the Shapley value, and has a number of basic economic motivations. We show that the price of stability for network design with respect to this fair cost allocation is $O(\log k)$, where k is the number of users, and that a good Nash equilibrium can be achieved via best-response dynamics in which users iteratively defect from a starting solution. This establishes that the fair cost allocation protocol is in fact a useful mechanism for inducing strategic behavior to form near-optimal equilibria. We discuss connections to the class of potential games defined by Monderer and Shapley, and extend our results to cases in which users are seeking to balance network design costs with latencies in the constructed network, with stronger results when the network has only delays and no construction costs. We also present bounds on the convergence time of best-response dynamics, and discuss extensions to a weighted game.

1. Introduction

In many network settings, the system behavior arises from the actions of a large number of independent agents, each motivated by self-interest and optimizing an individual objective function. As a result, the global performance of the system may not be as good as in a case where a central authority can simply dictate a solution; rather, we need to understand the quality of solutions that are consistent with self-interested behavior. Recent theoretical work has framed this type of question in the following general form: how much worse is the solution quality at a Nash equilibrium¹, relative to

¹Recall that a Nash equilibrium is a state of the system in which no agent has an interest in unilaterally changing its own behavior.

*Cornell University, Department of Computer Science, Upson Hall, Ithaca, NY 14853. Supported by ITR grant 0311333. Email: eanshel@cs.cornell.edu.

†Cornell University, Department of Computer Science, Upson Hall, Ithaca, NY 14853. Supported by the Department of Computer Science. Email: adg@cs.cornell.edu.

‡Cornell University, Department of Computer Science, Upson Hall, Ithaca, NY 14853. Email: kleinber@cs.cornell.edu. Supported in part by a David and Lucile Packard Foundation Fellowship and NSF grants 0081334 and 0311333.

§Cornell University, Department of Computer Science, Upson Hall, Ithaca, NY 14853. Supported in part by NSF grant CCR-032553, ITR grant 0311333, and ONR grant N00014-98-1-0589. Email: eva@cs.cornell.edu.

¶Cornell University, Department of Computer Science, Upson Hall, Ithaca, NY 14853. Supported by ITR grant 0311333. Email: wexler@cs.cornell.edu.

||UC Berkeley, Computer Science Division, Soda Hall, Berkeley, CA 94720. Supported by an NSF Postdoctoral Fellowship. Email: timr@cs.berkeley.edu.

the quality at a centrally enforced optimum? Questions of this genre have received considerable attention in recent years, for problems including routing [24, 25, 4], load balancing [5, 6, 16, 23], and facility location [26].

An important issue to explore in this area is the middle ground between centrally enforced solutions and completely unregulated anarchy. In most networking applications, it is not the case that agents are completely unrestricted; rather, they interact with an underlying protocol that essentially proposes a collective solution to all participants, who can each either accept it or defect from it. As a result, it is in the interest of the protocol designer to seek the *best* Nash equilibrium; this can naturally be viewed as the optimum subject to the constraint that the solution be *stable*, with no agent having an incentive to unilaterally defect from it once it is offered. Hence, one can view the ratio of the solution quality at the best Nash equilibrium relative to the global optimum as a *price of stability*, since it captures the problem of optimization subject to this constraint. Some recent work [1, 4] has considered this definition (termed the “optimistic price of anarchy” in [1]); it stands in contrast to the larger line of work in algorithmic game theory on the *price of anarchy* [21] — the ratio of the *worst* Nash equilibrium to the optimum — which is more suited to worst-case analysis of situations with essentially no protocol mediating interactions among the agents. Indeed, one can view the activity of a protocol designer seeking a good Nash equilibrium as being aligned with the general goals of mechanism design — producing a game that yields good outcomes when players act in their own self-interest.

Network Design Games. Network design is a natural area in which to explore the price of stability, given the large body of work in the networking literature on methods for sharing the cost of a designed network — often a virtual overlay, multicast tree, or other sub-network of the Internet — among a collection of participants. (See e.g. [9, 11] for overviews of work in this area).

A cost-sharing mechanism can be viewed as the underlying protocol that determines how much a network serving several participants will cost to each of them. Specifically, say that each user i has a pair of nodes (s_i, t_i) that it wishes to connect; it chooses an s_i - t_i path S_i ; and the cost-sharing mechanism then charges user i a cost of $C_i(S_1, \dots, S_k)$.

(Note that this cost can depend on the choices of the other users as well.) Although there are in principle many possible cost-sharing mechanisms, research in this area has converged on a few mechanisms with good theoretical and empirical behavior; here we focus on the following particularly natural one: the cost of each edge is shared equally by the set of all users whose paths contain it, so that $C_i(S_1, S_2, \dots, S_k) = \sum_{e \in S_i} \frac{c_e}{|\{j : e \in S_j\}|}$. This

equal-division mechanism has a number of basic economic motivations; it can be derived from the Shapley value [20], and it can be shown to be the unique cost-sharing scheme satisfying a number of different sets of axioms [9, 11, 20]. For the former reason, we will refer to it as the *Shapley cost-sharing mechanism*. Note that the total edge cost of the designed network is equal to the sum of the costs in the union of all S_i , and the costs allocated to users in the Shapley mechanism completely pay for this total edge cost: $\sum_{i=1}^n C_i(S_1, S_2, \dots, S_k) = \sum_{e \in \cup_i S_i} c_e$.

Now, the general question is to determine how this basic cost-sharing mechanism serves to influence the strategic behavior of the users, and what effect this has on the structure and overall cost of the network one obtains. Given a solution to the network design problem consisting of a vector of paths (S_1, \dots, S_k) for the n users, user i would be interested in deviating from this solution if there were an alternate s_i - t_i path S'_i so that changing to S'_i would lower its cost under the resulting allocation: $C_i(S_1, \dots, S_{i-1}, S'_i, S_{i+1}, \dots, S_k) < C_i(S_1, \dots, S_{i-1}, S_i, S_{i+1}, \dots, S_k)$. We say that a set of paths is a *Nash equilibrium* if no user has an interest in deviating. As we will see below, there exists a set of paths in Nash equilibrium for every instance of this network design game. (In this paper, we will only be concerned with *pure* Nash equilibrium; i.e. with equilibria where each user deterministically chooses a single path.)

The goal of a network design protocol is to suggest for each user i a path S_i so that the resulting set of paths is in Nash equilibrium and its total cost exceeds that of an optimal set of paths by as small a factor as possible; this factor is the *price of stability* of the instance. It is useful at this point to consider a simple example that illustrates how the price of stability can grow to a super-constant value (with k). Suppose k players wish to connect the common source s to their terminal t_i , assume player i has its own path of cost $1/i$, and all players can share a




Figure 1. An instance in which the price of stability converges to $H(k) = \Theta(\log k)$ as $\epsilon \rightarrow 0$.

common path of cost $1 + \epsilon$ for some small $\epsilon > 0$ (see Figure 1). The optimal solution would connect all agents through the common path for a total cost of $1 + \epsilon$. However, if this solution were offered to the users, they would defect from it one by one to their alternate paths. The unique Nash equilibrium has a cost of $\sum_{i=1}^k \frac{1}{i} = H(k)$.

While the price of stability in this instance grows with k , it only does so logarithmically. It is thus natural to ask how bad the price of stability can be for this network design problem. If we think about the example in Figure 1 further, it is also interesting to note that a good Nash equilibrium is reached by iterated greedy updating of players' solutions (in other words, best-response dynamics) starting from an optimal solution; it is natural to ask to what extent this holds in general.

Our Results. Our first main result is that in every instance of the network design problem with Shapley cost-sharing, there always exists a Nash equilibrium of total cost at most $H(k)$ times optimal. In other words, the simple example in Figure 1 is in fact the worst possible case.

We prove this result using a *potential function* method due to Monderer and Shapley [19] and Rosenthal [22] (see also [3]): one defines a potential function Φ on possible solutions and shows that any improving move by one of the users (i.e. to lower its own cost) reduces the value of Φ . Since the set of possible solutions is finite, it follows that any sequence of improving moves leads to a Nash equilibrium. The goal of Monderer and Shapley's and Rosenthal's work was to prove existence statements of this sort; for our purposes, we make further use of the potential function to prove a bound on

the price of stability. Specifically, we give bounds relating the value of the potential for a given solution to the overall cost of that solution; if we then iterate best-response dynamics starting from an optimal solution, the potential does not increase, and hence we can bound the cost of any solution that we reach. Thus, for this network design game, best-response dynamics starting from the optimum does in fact always lead to a good Nash equilibrium.

We can extend our basic result to a number of more general settings. To begin with, the $H(k)$ bound on the price of stability extends directly to the case in which users are selecting arbitrary subsets of a ground set (with elements' costs shared according to the Shapley value), rather than paths in a graph; it also extends to the case in which the cost of each edge is a non-decreasing concave function of the number of users on it. In addition, our results also hold if we introduce capacities into our model; each edge e may be used by at most u_e players, where u_e is the capacity of e .

We arrive at a more technically involved set of extensions if we wish to add latencies to the network design problem. Here each edge has a concave *construction cost* $c_e(x)$ when there are x users on the edge, and a *latency cost* $d_e(x)$; the cost experienced by a user is the full latency plus a fair share of the construction cost, $d_e(x) + c_e(x)/x$. We give general conditions on the latency functions that allow us to bound the price of stability in this case at $d \cdot H(k)$, where d depends on the delay functions used. Moreover, we obtain stronger bounds in the case where users experience only delays, not construction costs; this includes a result that relates the cost at the best Nash equilibrium to that of an optimum with twice as many players, and a result that improves

the potential-based bound on the price of stability for the single-source delay-only case.

Since a number of our proofs are obtained by following the results of best-response dynamics via a potential function, it is natural to investigate the speed of convergence of best-response dynamics for this game. We show that it converges to a Nash equilibrium in polynomial time for the case of two players, but that with k players, it can run for a time exponential in k . Whether there is a way to schedule players' moves to make best-response converge in a polynomial number of steps for this game in general is an interesting open question.

Finally, we consider a natural generalization of the cost-sharing model that carries us beyond the potential-function framework and raises interesting questions for further work. Specifically, suppose each user has a *weight* (perhaps corresponding to the amount of traffic it plans to send), and we change the cost-allocation so that user i 's payment for edge e is equal to the ratio of its weight to the total weight of all users on e . In addition to being intuitively natural, this definition is analogous to certain natural generalizations of the Shapley value [18]. The weighted model, however, is significantly more complicated: there is no longer a potential function whose value tracks improvements in users' costs when they greedily update their solutions, and it is an open question whether best-response dynamics will always converge to a Nash equilibrium. We have obtained some initial results here, including the convergence of best-response dynamics when all users seek to construct a path from a node s to a node t (the price of stability here is 1), and in the general model of users selecting sets from a ground set, when each element appears in the sets of at most two users. An interesting open question is to obtain more general results in this weighted setting, which appears to pose a challenge to potential-based methods. Further, we know that some results will necessarily look quite different in the weighted case; for example, using a construction involving user weights that grow exponentially in k , we can show that the price of stability can be as high as $\Omega(k)$.

Related Work. Network design games under a different model were considered by a subset of the authors in [1]; there, the setting was much more “unregulated” in that users could offer to pay for an arbitrary fraction of any edge in the network. This model resulted in instances where no pure Nash

equilibrium existed; and in many cases in [1] when pure Nash equilibria did exist, certain users were able to act as “free riders,” paying very little or nothing at all. The present model, on the other hand, ensures that there is always a pure Nash equilibrium within a logarithmic factor of optimal, in which users pay a fair portion of the resources they use. Network creation games of a fairly different flavor — in which users correspond to nodes, and can build subsets of the edges incident to them — have been considered in [2, 7, 10]. The model in this paper associates users instead with connection requests, and allows them to contribute to the cost of any edge that helps form a path that they need.

The bulk of the work on cost-sharing (see e.g. [9, 11] and the references there) tends to assume a fixed underlying set of edges. Jain and Vazirani [12] and Kent and Skorin-Kapov [15] consider cost-sharing for a single source network design game. Cost-sharing games assume that there is a central authority that designs and maintains the network, and decides appropriate cost-shares for each agent, depending on the graph and all other agents, via a complex algorithm. The agents' only role is to report their utility for being included in the network.

Here, on the other hand, we consider a simple cost-sharing mechanism, the Shapley-value, and ask what the strategic implications of a given cost-sharing mechanism are for the way in which a network will be designed. This question explores the feedback between the protocol that governs network construction and the behavior of self-interested agents that interact with this protocol. An approach of a similar style, though in a different setting related to routing, was pursued by Johari and Tsitsiklis [13]; there, they assumed a network protocol that priced traffic according to a scheme due to Kelly [14], and asked how this protocol would affect the strategic decisions of self-interested agents routing connections in the network.

The special case of our game with only delays is closely related of the congestion game of [25, 24]. They consider a game where the amount of flow carried by an individual user is infinitesimally small (a *non-atomic game*), while in this paper we assume that each user has a unit of flow, which it needs to route on a single path. In the non-atomic game of [25, 24] the Nash equilibrium is essentially unique (hence there is no distinction between the price of anarchy and stability), while in our atomic game there can be many equilibria. Fabrikant, Papadim-

itriou, and Talwar [8] consider our atomic game with delays only. They give a polynomial time algorithm to minimize the potential function Φ in the case that all users share a common source, and show that finding any equilibrium solution is PLS-complete for multiple source-sink pairs. Our results extend the price of anarchy results of [25, 24] about non-atomic games to results on the price of stability for the case of single source atomic games.

A weighted game similar to our is presented by Libman and Orda [17], with a different mechanism for distributing costs among users. They do not consider the price of stability, and instead focus on convergence in parallel networks.

2. Nash Equilibria of Network Design with Shapley Cost-Sharing

In this section we consider the *Fair Connection Game* for k players as defined in the Introduction. Let a directed graph $G = (V, E)$ be given, with each edge having a nonnegative cost c_e . Each player i has a set of terminal nodes T_i that he wants to connect. A strategy of a player is a set of edges $S_i \subset E$ such that S_i connects all nodes in T_i . We assume that we use the Shapley value to share the cost of the edges, i.e. all players using an edge split up the cost of the edge equally. Given a vector of players' strategies $S = (S_1, \dots, S_k)$, let x_e be the number of agents whose strategy contains edge e . Then the cost to agent i is $C_i(S) = \sum_{e \in S_i} (c_e/x_e)$, and the goal of each agent is to connect its terminals with minimum total cost.

In the worst case, Nash equilibria can be very expensive in this game, so that the price of anarchy becomes as large as k . To see this, consider k players with common source s and sink t , and two parallel edges of cost 1 and k . The worst equilibrium has all players selecting the more expensive edge, thereby paying k times the cost of the optimal network. However, we can bound the price of stability by $H(k)$, which is the harmonic sum $1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{k}$, as follows.

Theorem 2.1 *The price of stability of the fair connection game is at most $H(k)$.*

Proof: The fair connection game that we have defined falls into the class of congestion games as defined by Monderer and Shapley [19], as the cost of an edge e to a user i is $f_e(x) = c_e/x$, which depends

only on edge e and the number of users x whose strategy contains e . Monderer and Shapley [19] show that all congestion games have deterministic Nash equilibria. They prove this using a potential function Φ , defined as follows.

$$\Phi(S) = \sum_{e \in E} \sum_{x=1}^{x_e} f_e(x) \quad (1)$$

Monderer and Shapley [19] show that for any strategy $S = (S_1, \dots, S_k)$ if a single player i deviates to strategy S'_i , then the change in the potential value $\Phi(S) - \Phi(S')$ of the new strategy set $S' = (S_1, \dots, S'_i, \dots, S_k)$ is exactly the change in the cost to player i . Note that the change of player i 's strategy affects the cost of many other players $j \neq i$, but the Φ value is not effected by the change in the cost of these players, it simply tracks the cost of the player who changes its strategy. They call a game in which such a function Φ exists a *potential game*. To show that such a potential game has a deterministic Nash equilibrium, start from any state $S = (S_1, \dots, S_k)$ and consider a sequence of selfish moves (allowing players to change strategies to improve their costs). In a congestion game any sequence of such improving moves leads to a Nash equilibrium as each move decreases the potential function Φ , and hence must lead to a stable state.

Monderer and Shapley do not say anything about the quality of Nash equilibria with respect to the centralized optimum, but we can use their potential function to establish our bound. Let x_e be defined as above with respect to S . Now the potential function of Equation 1 in our case is $\Phi(S) = \sum_{e \in E} c_e H(x_e)$. According to the above argument, any improving deviation decreases $\Phi(S)$, and so a sequence of improving deviations by players must eventually result in a Nash equilibrium.

Consider the strategy $S^* = (S_1^*, \dots, S_k^*)$ defining the optimal centralized solution. Let $OPT = \sum_{e \in S^*} c_e$ be the cost of this solution. Then, $\Phi(S^*) \leq \sum_{e \in S^*} (c_e \cdot H(k))$, which is exactly $H(k) \cdot OPT$. Now we start from strategy S^* and follow a sequence of improving self-interested moves. We know that this will result in a Nash equilibrium S with $\Phi(S) \leq \Phi(S^*)$.

Note that the potential value of any solution S is at least the total cost: $\Phi(S) \geq \sum_{e \in S} c_e = cost(S)$. Therefore, there exists a Nash equilibrium with cost at most $H(k) \cdot OPT$, as desired. ■

Recall from the Introduction that this bound is

tight as shown by the example in Figure 1. Unfortunately, even though Theorem 2.1 says that cheap Nash equilibria exist, finding them is NP-complete (by a reduction from 3D-Matching).

We can extend the results of Theorem 2.1 to concave cost functions. Consider the extended fair connection game where instead of a constant cost c_e , each edge has a cost which depends on the number of players using that edge, $c_e(x)$. We assume that $c_e(x)$ is a nondecreasing, concave function, modeling the buy-at-bulk economy of scale of buying edges that can be used by more players. Notice that the cost of an edge $c_e(x)$ might increase with the number of players using it, but the cost per player $f_e(x) = c_e(x)/x$ decreases if $c_e(x)$ is concave.

Theorem 2.2 *Take a fair connection game with each edge having a nondecreasing concave cost function $c_e(x)$, where x is the number of players using edge e . Then the price of stability is at most $H(k)$.*

Proof: The proof is analogous to the proof of Theorem 2.1. We use the potential function $\Phi(S)$ defined by (1). As before, the change in potential if a player i deviates equals exactly to the change of that player's payments. We start with the strategy S^* with minimum total cost, and perform a series of improving deviations until we reach a Nash equilibrium S with $\Phi(S) \leq \Phi(S^*)$. To finish the proof all we need to show is that $cost(S) \leq \Phi(S) \leq H(k) \cdot cost(S)$ for all strategies S . The second inequality follows since $c_e(x)$ is nondecreasing and therefore $\sum_{x=1}^{x_e} (c_e(x)/x) \leq H(x_e) \cdot c_e(x_e)$. To see that $cost(S) \leq \Phi(S)$ notice that since $c_e(x)$ is concave, the cost per player must decrease with x , i.e. $c_e(x)/x$ is a nonincreasing function. Therefore, $cost(S) = \sum_{e \in S} c_e(x_e) = \sum_{e \in S} x_e \cdot (c_e(x_e)/x_e) \leq \Phi(S)$, which finishes the proof. ■

Extensions. The proof of Theorem 2.2 extends to a general congestion game, where players attempt to share a set of resources R they need. Instead of having an underlying graph structure, we now think of each $s \in R$ as a resource with a concave cost function $c_s(x)$ of the number of users selecting sets containing s . The possible strategies of each player i is a set \mathcal{S}_i of subsets of R . Each player seeks to select a set $S_i \in \mathcal{S}_i$ so as to minimize his cost. Since the proofs above did not rely on the graph structure, they translate directly to this extension.

We can further extend the results to the case when the cost to a player is a combination of the cost

$c_e(x)/x$, and a function of the selected set, such as the distance between terminals in the network design case. More precisely, the price of stability is still at most $H(k)$ if each player is trying to minimize the cost $\sum_{e \in S_i} (c_e(x_e)/x_e) + d_i(S_i)$ where c_e is monotone increasing and concave, and d_i is an arbitrary function specific to player i (e.g. a distance function, or diameter of S_i , etc.). The proof is analogous to Theorem 2.2, except with a new potential $\Phi(S) = \sum_i d_i(S_i) + \sum_{e \in S} \sum_{x=1}^{x_e} \frac{c_e(x)}{x}$. Notice that this is technically not a congestion game on the given graph G . Finally we note that all these results (as well as those subsequent) hold in the presence of capacities. Adding capacities u_e to each edge e and disallowing more than u_e players to use e at any time does not substantially alter any of our proofs.

The Case of Undirected Graphs. While the bound of $H(k)$ is tight for general directed graphs, it is not tight for undirected graphs. Finding the correct bound is an interesting open problem. In the case of two players, our bound on the price of stability is $H(2) = 3/2$. In the full version we show that that this bound can be improved to $4/3$ in the case of two players and a single source. We also give an example to show that this bound is tight.

3. Dealing with Delays

In most of the previous section, we assumed that the utility of a player depends only on the cost of the edges he uses. What changes if we introduce latency into the picture? We have extended this to the case when the players' cost is a combination of "design" cost and the length of the path selected. More generally, delay on an edge does not have to be simply the "hop-count", but can also depend on congestion, i.e., on the number of players using the edge. In this section we will consider such a model.

Assume that each edge has both a cost function $c_e(x)$ and a latency function $d_e(x)$, where $c_e(x)$ is the cost of building the edge e for x users and the users will share this cost equally, while $d_e(x)$ is the delay suffered by users on edge e if x users are sharing the edge. The goal of each user will be to minimize the sum of his cost and his latency. If we assume that both the cost and latency for each edge depend only on the number of players using that edge, then this fits directly into our model of a congestion game above: the total cost felt by each user on the edge is $f_e(x) = c_e(x)/x + d_e(x)$. If the func-

tion $xf_e(x)$ is concave then Theorem 2.2 applies. But while concave functions are natural for modeling cost, latency tends to be convex.

3.1. Combining Costs and Delays

First, we extend the argument in the proof of Theorem 2.2 to general functions f_e . The most general version of this argument is expressed in the following theorem.

Theorem 3.1 *Consider a fair connection game with arbitrary edge-cost functions f_e . Suppose that $\Phi(S)$ is as in Equation 1, with $\text{cost}(S) \leq A \cdot \Phi(S)$, and $\Phi(S) \leq B \cdot \text{cost}(S)$ for all S . Then, the price of stability is at most $A \cdot B$.*

Proof: Let S^* be a strategy such that S_i^* is the set of edges i uses in the centralized optimal solution. We know from above that if we perform a series of improving deviations on it, we must converge to a Nash equilibrium S' with potential value at most $\Phi(S^*)$. By our assumptions, $\text{cost}(S') \leq A \cdot \Phi(S') \leq A \cdot \Phi(S^*) \leq AB \cdot \text{cost}(S^*) = AB \cdot \text{OPT}$. ■

Our main interest in this section are functions $f_e(x)$ that are the sums of the fair share of a cost and a delay, i.e., $f_e(x) = c_e(x)/x + d_e(x)$. We will assume that $d_e(x)$ is monotone increasing, while $c_e(x)$ is monotone increasing and concave.

Corollary 3.2 *If $c_e(x)$ is concave and nondecreasing, $d_e(x)$ is nondecreasing for all e , and $x_e d_e(x_e) \leq A \sum_{x=1}^{x_e} d_e(x)$ for all e and x_e , then the price of stability is at most $A \cdot H(k)$. In particular, if $d_e(x)$ is a polynomial with degree at most l and nonnegative coefficients, then the price of stability is at most $(l+1) \cdot H(k)$.*

Proof: For functions $f_e(x) = c_e(x)/x + d_e(x)$, both the cost and potential of a solution come in two parts corresponding to the cost c and delay d .

For the part corresponding to cost the potential over-estimates the cost by at most a factor of $H(k)$ as proved in Theorem 2.2. If on the delay, the potential underestimates the cost by at most a factor of A , then we get the bound of $A \cdot H(k)$ for the price of stability by Theorem 3.1. ■

Therefore, for reasonable delay functions, the price of stability cannot be too large. In particular, if the utility function of each player depends on

a concave cost and delay that is independent of the number of users on the edge, then we get that the price of stability is at most $H(k)$ as we have shown at the end of the previous section. If the delay grows linearly with the number of users, then the price of stability is at most $2H(k)$.

3.2. Games with Only Delays

In this subsection we consider games with only delay. We assume that the cost of a player for using an edge e used by x players is $f_e(x) = d_e(x)$, and d_e is a monotone increasing function of x . This cost function models delays that are increasing with congestion.

We will consider the special case when there is a common source s . Each player i has one additional terminal t_i , and the player wants to connect s to t_i via a directed path. Fabrikant, Papadimitriou, and Talwar [8] showed that in this case, one can compute the Nash equilibrium minimizing the potential function Φ via a minimum cost flow computation. For each edge e they introduce many parallel copies, each with capacity 1, and cost $d_e(x)$ for integers $x > 0$. We will use properties of a minimum cost flow for establishing our results.

First we show a bicriteria bound, and compare the cost of the cheapest Nash equilibrium to that of the optimum design with twice as many players.

Theorem 3.3 *Consider the single source case of a congestion game with only delays. Let S be the minimum cost Nash equilibrium and S^* be the minimum cost solution for the problem where each player i is replaced by two players. Then $\text{cost}(S) \leq \text{cost}(S^*)$.*

Proof: Consider the Nash equilibrium obtained by Fabrikant et al [8] via a minimum cost flow computation. Assume that x_e is the number of users using edge e at this equilibrium. By assumption, all users share a common source s . Let $D(v)$ denote the cost of the minimum cost path in the residual graph from s to v . The length of the path of user i is at most $D(t_i)$ (as otherwise the residual graph would have a negative cycle) and hence we get that $\text{cost}(S) \leq \sum_i D(t_i)$.

Now consider a modified delay function \hat{d}_e for each edge $e = (u, v)$. Define $\hat{d}_e(x) = d_e(x)$ if $x > x_e$, and $\hat{d}_e(x) = D(v) - D(u)$ if $x \leq x_e$. Note that for any edge e we have $D(v) - D(u) \leq d_e(x_e + 1)$ as the edge $e = (u, v)$ is in the residual graph with cost

$d_e(x_e + 1)$. This implies that the modified delay \hat{d} is monotone. For edges with $x_e \neq 0$ we also have that $d_e(x_e) \leq D(v) - D(u)$ as the reverse edge (v, u) is in the residual graph with cost $-d_e(x_e)$, so the delay of an edge is not decreased.

Now observe that, subject to the new delay \hat{d} , the shortest path from s to t_i is length $D(t_i)$ even in an empty network. The minimum possible cost of two paths from s to t_i for the two users corresponding to user i is then at least $2D(t_i)$ for each player i . Therefore the minimum cost of a solution with delays \hat{d} is at least $2 \sum_i D(t_i)$.

To bound $\text{cost}(S^*)$ we need to bound the difference in cost of a solution when measured with delays \hat{d} and d . Note that for any edge $e = (u, v)$ and any number x we have that $x\hat{d}_e(x) - xd_e(x) \leq x_e(D(v) - D(u))$, and hence the difference in total cost is at most $\sum_{e=(u,v)} x_e(D(v) - D(u)) = \sum_i D(t_i)$. Using this, we get that $\text{cost}(S^*) \geq \sum_i D(t_i) \geq \text{cost}(S)$. ■

Note that a similar bound is not possible for a model with both costs and delays, when additional users compensate to some extent for the price of stability. Consider a problem with two parallel links e and e' and k users. Assume on link e the cost is all design cost $c_e(x) = 1 + \varepsilon$ for a small $\varepsilon > 0$. On the other link e' the cost is all delay, and the delay with x users is $d_{e'}(x) = 1/(k - x + 1)$. The optimum solution is to use the first edge e , and it costs $1 + \varepsilon$. Note that the optimum with any number of extra users costs the same, as this is all design cost. On the other hand, the only Nash is to have all users on e' , incurring delay 1, for a total cost of k .

Note that the $H(k)$ term in Corollary 3.2 comes from the concave cost c , and so the bound obtained there improves by an $H(k)$ factor when the cost consists of only delay. Roughgarden [24] showed a tighter bound for non-atomic games. He assumed that the delay is monotone increasing, and the total cost of an edge $xd_e(x)$ is a convex function of the traffic x . He showed that for any class of such functions \mathcal{D} containing all constant functions, the price of anarchy is always obtained on a two node, two link network. Let us call $\alpha(\mathcal{D})$ the price of anarchy for non-atomic games with delays from the class \mathcal{D} (which is also the price of stability, since the Nash equilibrium is unique). For example, Roughgarden [24] showed that for polynomials of degree at most l this bound is $O(l/\log l)$. Here we extend this result to a single source atomic game.

Theorem 3.4 *If in a single source fair connection game all costs are delays, and all delays are from a set \mathcal{D} satisfying the above condition, then the price of stability is at most $\alpha(\mathcal{D})$.*

Proof Sketch: We defer the full proof to the extended version. The idea is as follows. We construct a modified network \hat{G} by adding edges and capacities to G . We show that the Nash equilibrium is not affected by the change, and the optimum can only improve. We obtain the claimed bound by comparing the cost of the Nash equilibrium to the minimum cost of a fractional solution (a flow) in \hat{G} .

Consider the Nash equilibrium obtained via a minimum cost flow computation as in the proof of Theorem 3.3, let x_e be the number of paths using edge e , and $D(v)$ be the length of the shortest path from s to v in the residual graph. Add to each edge $e = (u, v)$ a capacity of x_e , and augment our network by adding a parallel edge e' with constant delay $D(v) - D(u)$. Note that the new capacity and the added links do not effect the equilibrium. We will show that for each edge e , the two parallel copies: edge e with new capacity x_e and edge e' , can carry any number of paths at least as cheaply as the original edge e could. Hence this change in the network only improves the minimum possible cost. Let \hat{G} denote the resulting network flow problem.

We will show that the minimum cost fractional flow in \hat{G} is obtained by splitting the flow x_e between the two edges e and e' appropriately to make the cost of the gradient equal. The claimed bound will then follow by comparing the cost $x_e d_e(x_e)$ of the edge at Nash equilibrium with the cost of the corresponding two edges e and e' in \hat{G} . ■

4. Convergence of Best Response

In this section, we address the convergence properties of best response dynamics in our game.

Theorem 4.1 *In the two player fair connection game, best response dynamics starting from any configuration converges to a Nash equilibrium in polynomial time.*

The detailed proof appears in the extended version and shows that for any best response run, the number of edges shared by both players increases monotonically. For more players, however, the hope of any positive result about best response dynamics seems slim. In fact, we can show the following.

Theorem 4.2 *Best response dynamics for k players may run in time exponential in k .*

The proof constructs an example of a game that can simulate a k -bit counter. The extended version contains details of the construction.

5. Weighted Players

So far we have assumed that players sharing an edge e pay equal fractions of e 's cost. We now consider a game with fixed edge costs where players have weights $w_i \geq 1$, and players' payments are proportional to their weight. More precisely, given a strategy $S = (S_1, \dots, S_k)$, define W to be the total weight of all players, and let W_e be the sum of the weights of players using e . Then player i 's payment for edge e will be $\frac{w_i}{W_e} c_e$.

Note that the potential function $\Phi(S)$ used for the unweighted version of the game is not a potential function once weights are added. In particular, in a weighted game, improving moves can increase the value of $\Phi(S)$, as this is no longer a congestion game. The following theorem uses a new potential function for a special class of weighted games.

Theorem 5.1 *In a weighted game where each edge e is in the strategy spaces of at most two players, there exists a potential function for this game, and hence a Nash equilibrium exists.*

Proof: Consider the following potential function. For each edge e used by players i and j , define

$$\Phi_e(S) = \begin{cases} c_e w_i & \text{if player } i \text{ uses } e \text{ in } S \\ c_e w_j & \text{if player } j \text{ uses } e \text{ in } S \\ c_e \theta_{ij} & \text{if both players } i \text{ and } j \text{ use } e \text{ in } S \\ 0 & \text{otherwise} \end{cases}$$

where $\theta_{ij} = (w_i + w_j - \frac{w_i w_j}{w_i + w_j})$. For any edge e with only one player i , simply set $\Phi_e(S) = w_i c_e$ if i uses e and 0 otherwise. Define $\Phi(S) = \sum_e \Phi_e(S)$. We now simply need to argue that if a player makes an improving move, then $\Phi(S)$ decreases. Consider a player i and an edge e that player i joins. If the edge already supported another player j , then i 's cost for using e is $c_e \frac{w_i}{w_i + w_j}$, while the change in $\Phi_e(S)$ is

$$c_e \left(w_i - \frac{w_i w_j}{w_i + w_j} \right) = c_e \frac{w_i^2}{w_i + w_j}.$$

Thus the change in potential when i joins e equals the cost i incurs, scaled up by a factor of w_i . In

fact, it is easy to show the more general fact that when player i moves, the change in $\Phi(S)$ is equal to the change in player i 's payments scaled up by w_i . This means that improving moves always decrease $\Phi(S)$, thus proving the theorem. ■

Note that this applies not only to paths, but also to the generalized model in which players select subsets from some ground set. The analogous condition is that no ground element appears in the strategy spaces of more than two players.

Corollary 5.2 *Any two-player weighted game has a Nash equilibrium.*

While the above potential function also implies a bound on the price of stability, even with only two players this bound is very weak. However, if there are only two players with weights 1 and $w \geq 1$, then we can show that the price of stability is at most $1 + \frac{1}{1+w}$, and this is tight for all w .

The following result shows the existence of Nash equilibria in weighted single commodity games.

Theorem 5.3 *For any weighted game in which all players have the same source s and sink t , best response dynamics converges to a Nash equilibrium, and hence Nash equilibria exist.*

Proof: Start with any initial set of strategies S . For every $s - t$ path P define the marginal cost of P to be $c(P) = \sum_{e \in P} \frac{c_e}{W_e}$ where W_e depends on S . Observe that if player i currently uses path P , then i 's payment is $w_i c(P)$. Define $P(S)$ to be a tuple of the values $c(P)$ over all paths P , sorted in increasing order. We want to show that the cheapest improving deviation of any player causes $P(S)$ to strictly decrease lexicographically.

Suppose that one of the best moves for player i is to switch paths from P_1 to P_2 . Let \mathcal{P} denote the set of paths that intersect $P_1 \cup P_2$. For any pair of paths P and Q , let $c_P(Q)$ denote the new value of $c(Q)$ after player i has switched to path P . To show that $P(S)$ strictly decreases lexicographically, it suffices to show that

$$\min_{P \in \mathcal{P}} c_{P_2}(P) < \min_{P \in \mathcal{P}} c(P). \quad (2)$$

Define $P' = \arg \min_{P \in \mathcal{P}} c(P)$. Since P_2 was i 's best response, $c_{P_2}(P_2) \leq c_P(P)$ for all paths P . In particular, $c_{P_2}(P_2) \leq c_{P'}(P')$. We also know that $c_{P'}(P') \leq c(P')$, since in deviating to P' , player i adds itself to some edges of P' . In fact,

$c_{P'}(P') < c(P')$ unless $P' = P_1$. Assuming $P' \neq P_1$, we now have that $c_{P_2}(P_2) < c(P')$, which proves inequality 2. If $P' = P_1$, then since player i decided to deviate, $c_{P_2}(P_2) < c(P_1)$. Therefore, we once again have that $c_{P_2}(P_2) < c(P')$, as desired. ■

In the case where the graph consists of only 2 nodes s and t joined by parallel links, we can similarly show that any sequence of improving responses converge to a Nash equilibrium.

With arbitrarily increasing cost functions, [17] gives an example demonstrating that a weighted game may not have any pure Nash equilibria. However, it is still an open problem to determine whether weighted games with fixed costs always have Nash Equilibria. While the authors believe they do, it is not clear how to adapt a potential-style argument to prove this. The construction above does not even extend to games where 3 players share an edge. However, in either case, the following claim shows that the price of stability bounds from the unweighted case will not carry over.

Theorem 5.4 *There are weighted games for which the price of stability is $\Theta(\log W)$ and $\Theta(k)$.*

An example exhibiting this is a modified version of the graph in Figure 1. Change the edge with cost $1 + \varepsilon$ to cost 1, and for all other edges with positive cost, set the new cost to be $\frac{1}{2}$. For $1 \leq i \leq k$ let player i have weight $w_i = 2^{i-1}$. Since each player has a greater weight than all smaller weight players combined, the only Nash equilibrium has cost $\frac{k}{2} = \Theta(\log W)$, while the optimal solution has cost 1.

References

- [1] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *STOC*, 2003.
- [2] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68, 2000.
- [3] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [4] J. Correa, A. Schulz, and N. S. Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 2003.
- [5] A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. *STOC*, 2002.
- [6] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *SODA*, 2002.
- [7] A. Fabrikant, A. Luthra, E. Maneva, S. Papadimitriou, and S. Shenker. On a network creation game. *PODC*, 2003.
- [8] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. *STOC*, 2004.
- [9] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *JCSS*, 63:21–41, 2001.
- [10] H. Heller and S. Sarangi. Nash networks with heterogeneous agents. *Working Paper Series*, E-2001-1, Virginia Tech, 2001.
- [11] S. Herzog, S. Shenker, and D. Estrin. Sharing the “cost” of multicast trees: An axiomatic analysis. *IEEE/ACM Transactions on Networking*, Dec. 1997.
- [12] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. *STOC*, 2001.
- [13] R. Johari and J. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, to appear.
- [14] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [15] K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on mst’s. *Operations Research Proceedings KOI*, pages 43–48, 1996.
- [16] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *STACS*, 1999.
- [17] L. Libman and A. Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17:4:385–409, 2001.
- [18] D. Monderer and D. Samet. *In Handbook of Game Theory Vol. III*. Elsevier Science, 2002.
- [19] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [20] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Duke Economics Working Paper*, 96-31, 1996.
- [21] C. Papadimitriou. Algorithms, games, and the internet. *STOC*, 2001.
- [22] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3:53–59, 1973.
- [23] T. Roughgarden. Stackelberg scheduling strategies. *STOC*, 2001.
- [24] T. Roughgarden. The price of anarchy is independent of the network topology. *STOC*, 2002.
- [25] T. Roughgarden and É. Tardos. How bad is selfish routing? *FOCS*, 2000.
- [26] A. Vetta. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. *FOCS*, 2002.