Network Design with Weighted Players

Ho-Lin Chen[*] Tim Roughgarden[†]

June 21, 2007

Abstract

We consider a model of game-theoretic network design initially studied by Anshelevich et al. [2], where selfish players select paths in a network to minimize their cost, which is prescribed by Shapley cost shares. If all players are identical, the cost share incurred by a player for an edge in its path is the fixed cost of the edge divided by the number of players using it. In this special case, Anshelevich et al. [2] proved that pure-strategy Nash equilibria always exist and that the price of stability—the ratio between the cost of the best Nash equilibrium and that of an optimal solution—is $\Theta(\log k)$, where k is the number of players. Little was known about the existence of equilibria or the price of stability in the general *weighted* version of the game. Here, each player i has a weight $w_i \geq 1$, and its cost share of an edge in its path equals w_i times the edge cost, divided by the total weight of the players using the edge.

This paper presents the first general results on weighted Shapley network design games. First, we give a simple example with no pure-strategy Nash equilibrium. This motivates considering the price of stability with respect to α-approximate Nash equilibria—outcomes from which no player can decrease its cost by more than an α multiplicative factor. Our first positive result is that $O(\log w_{max})$-approximate Nash equilibria exist in all weighted Shapley network design games, where w_{max} is the maximum player weight. More generally, we establish the following trade-off between the two objectives of good stability and low cost: for every $\alpha = \Omega(\log w_{max})$, the price of stability with respect to $O(\alpha)$-approximate Nash equilibria is $O((\log W)/\alpha)$, where W is the sum of the players' weights. In particular, there is always an $O(\log W)$-approximate Nash equilibrium with cost within a constant factor of optimal.

Finally, we show that this trade-off curve is nearly optimal: we construct a family of networks without $o(\log w_{max}/ \log\log w_{max})$-approximate Nash equilibria, and show that for all $\alpha = \Omega(\log w_{max}/\log\log w_{max})$, achieving a price of stability of $O(\log W/\alpha)$ requires relaxing equilibrium constraints by an $\Omega(\alpha)$ factor.

[*]Department of Computer Science, Stanford University, 393 Terman Engineering Building, Stanford CA 94305. Research supported in part by NSF Award 0323766. Email: `holin@stanford.edu`.

[†]Department of Computer Science, Stanford University, 462 Gates Building, Stanford CA 94305. Supported in part by by ONR grant N00014-04-1-0725, DARPA grant W911NF-04-9-0001, and an NSF CAREER Award. Email: `tim@cs.stanford.edu`.

1 Introduction

1.1 The Price of Stability in Network Design Games

Understanding the interaction between incentives and optimization in networks is an important problem that has recently been the focus of much work by the theoretical computer science community. Despite the wealth of results obtained in this area over the past five years, *network design and formation* remains a fundamental topic that is not well understood. While economists and social scientists have long studied game-theoretic models for how networks are or should be created with self-interested agents (see e.g. [6, 14, 15] and the references therein), the mathematical techniques for quantifying the performance of such networks are currently limited.

The goal of quantifying performance (or lack thereof) in the presence of selfish behavior naturally motivates the twin concepts of the *price of anarchy* and the *price of stability.* To define these, first recall that a *(pure-strategy) Nash equilibrium* is an assignment of all of the players of a noncooperative game to strategies so that the following stability property holds: no player can switch strategies and become better off, assuming that all other players hold their strategies fixed. As the outcome of selfish, uncoordinated behavior, Nash equilibria are typically inefficient and do not optimize natural objective functions [11].

The price of anarchy and the price of stability are two ways to measure the inefficiency of Nash equilibria of a game, with respect to a notion of "social good" (such as the total cost incurred by all of the players). The price of anarchy of a game, first defined in Koutsoupias and Papadimitriou [16], is the ratio of the objective function value of the *worst* Nash equilibrium and that of an optimal solution. The price of anarchy is natural from the perspective of worst-case analysis—an upper bound on the price of anarchy bounds the inefficiency of every possible stable outcome of a game.

The price of stability, by contrast, is the ratio of the objective function value of the *best* Nash equilibrium and that of an optimal solution. The price of stability was first studied in Schulz and Stier Moses [25] and was so-called in Anshelevich et al. [2]. The price of stability has primarily been studied in network design games [2, 3], with the interpretation that the network will be designed by a central authority (for use by selfish agents), but that this authority is unable or unwilling to prevent the network users from acting selfishly after the network is built. In such a setting, the best Nash equilibrium—the best network that accounts for the incentives facing the network users—is an obvious solution to propose. In this sense, the price of stability measures the necessary degradation in solution quality caused by imposing the game-theoretic constraint of stability.

1.2 Shapley Cost Sharing with Unweighted Players

The goal of analyzing the cost of networks created by or designed for selfish users was first proposed by Papadimitriou [21] and initially explored independently by Anshelevich et al. [3] and Fabrikant et al. [12]. These two papers studied different types of network design games; also, the first considered the price of stability (where it was called the "optimistic price of anarchy"), the second the price of anarchy. (See also [1, 10, 18] for more recent work on these and related models.) Closest to the present work is a variation on the model of [3] that was proposed and studied by Anshelevich et at. [2], which they called *network design with Shapley cost sharing* and we will abbreviate to *Shapley network design games*.

The most basic model considered in [2] is the following. The game occurs in a directed graph $G = (V, E)$, where each edge e has a nonnegative cost c_e, and each player i is identified with a source-sink pair (s_i, t_i). Every player i picks a path P_i from its source to its destination, thereby creating the network $(V, \cup_i P_i)$ at a social cost of $\sum_{e \in \cup_i P_i} c_e$. This social cost is shared among the players in the following way. First, if edge e lies in f_e of the chosen paths, then each player choosing such a path pays a proportional share $\pi_e = c_e/f_e$ of the cost. The overall cost $c_i(P_1, \ldots, P_k)$ to player i is then the sum $\sum_{e \in P_i} \pi_e$ of these proportional shares.

This proportional sharing method enjoys numerous desirable properties. It is "budget balanced", in that it partitions the social cost among the players; it can be derived from the Shapley value, and as a consequence is the unique cost-sharing method satisfying certain fairness axioms (see e.g. [19]); and, as shown in [2], it coaxes benign behavior from the players. Specifically, Anshelevich et al. [2] showed that a pure-strategy Nash equilibrium always exists—in contrast to the more general cost-sharing that was allowed in the predecessor model [3]—and that the price of stability under Shapley cost-sharing is at most the kth harmonic number $\mathcal{H}_k = O(\log k)$, where k is the number of players. Anshelevich et al. [2] also provided an example showing that this upper bound is the best possible, and proved numerous extensions.

1.3 Shapley Cost Sharing with Weighted Players

A natural and important extension that Anshelevich et al. [2] identified but proved few results for is that to *weighted* players. In most network design settings, we expect the amount of traffic to vary across source-sink pairs. Such non-uniformity arises for many reasons. For example, players could represent populations of customers of Internet Service Providers, which cannot be expected to possess a common size; players could represent individuals with different bandwidth requirements; or collusion among several players could yield a single "virtual" player with size equal to the sum of those of the colluding players.

The definition of network design with Shapley cost-sharing extends easily to include weighted players: if w_i denotes the weight of player i, then i's cost share of an edge e is $c_e \cdot w_i/W_e$, where W_e is the total weight of the players that use a path containing the edge e. While easy to define, this weighted network design game appeared challenging to analyze. Indeed, prior to the present work, the primary results known for this weighted game were essentially suggestions that it exhibits more complex behavior than its unweighted counterpart. In particular, Anshelevich et al. [2] proved the following: that the key "potential function" proof technique for the unweighted case cannot be directly used for games with weighted players; and that the price of stability can be as large as $\Omega(k + \log W)$, where k is the number of players and $W = \sum_i w_i$ is the sum of the players' weights (assuming $w_i \geq 1$ for all i). The positive results of [2] for weighted games concerned only the special cases of 2-player games and single-source, single-sink games, where pure-strategy Nash equilibria were shown to exist. No further positive or negative results on either the existence of pure-strategy Nash equilibria or on the price of stability were known for weighted Shapley network design games.

1.4 Our Results

In this paper, we give the first general results for weighted Shapley network design games. We set the stage for our work in Section 3 by exhibiting such a game with no pure-strategy Nash equilibrium. This example has only three players, employs a single-sink undirected network, and the ratio between the maximum and minimum player weights can be made arbitrarily small. (Pure-strategy Nash equilibria are known to exist in all weighted Shapley network design games with two players [2].) Thus there are no large classes of weighted Shapley network design games that always possess pure-strategy Nash equilibria beyond those identified in [2].

Our example motivates considering a larger class of equilibria to recover a guarantee that equilibria exist. Once existence has been established, we can then attempt to bound the price of stability with respect to this larger set of equilibria. There are several possible approaches to accomplishing this goal, and we compare these at length in the next subsection. In this paper, we pursue the same line of inquiry as in Anshelevich et al. [3]—where for a different but related network design game, pure-strategy Nash equilibria did not necessarily exist—and consider *approximate* pure-strategy Nash equilibria. An outcome is an *α-approximate Nash equilibrium* if no player can decrease its cost by more than an α multiplicative factor by deviating. The obvious goal is then to prove that α-approximate Nash equilibria always exist and that some such equilibrium has cost within a β factor of optimal, where α and β are as small as possible. Since these two parameters work against each other, we seek to understand more generally the interaction between the best-possible values of α and β. How much stability must we give up in order to achieve a low-cost solution, and vice versa? Is it possible to take one or both of α, β to be an absolute constant? The present paper is the first to study the trade-off curve for these two parameters in Shapley network design games.

Our main results give a complete solution to these questions. To describe them, scale players' weights so that the minimum player weight is 1, and let w_{max} and W denote the maximum weight and the sum of all weights, respectively. On the positive side, we show that every weighted Shapley network design game admits an $O(\log w_{max})$-approximate Nash equilibrium, and that the price of stability with respect to such equilibria is $O(\log W)$. More generally, we prove the following trade-off between the two objectives: for every $\alpha = \Omega(\log w_{max})$, the price of stability with respect to $O(\alpha)$-approximate Nash equilibria is $O((\log W)/\alpha)$. Thus to implement a network with cost within a constant factor of the optimal solution, it suffices to relax the equilibrium constraints by a logarithmic (in W) factor. This is a new result even for unweighted Shapley network design games. (Recall that in unweighted games, it is impossible to approximate the cost to within an $o(\log k)$ factor without relaxing the equilibrium constraints [2].)

On the negative side, we demonstrate that this trade-off curve is very close to the best possible. In our most involved construction, we exhibit a family of weighted Shapley network design games without $o(\log w_{max}/\log\log w_{max})$-approximate Nash equilibria. Recovering the existence of equilibria therefore requires relaxing the equilibrium constraints by a super-constant (though only logarithmic) factor. We also show that for every $\alpha = \Omega(\log w_{max}/\log\log w_{max})$, a price of stability of $O((\log W)/\alpha)$ can only be obtained by relaxing the equilibrium constraints by an $\Omega(\alpha)$ factor.

1.5 Discussion of Alternative Approaches

We conclude the Introduction by justifying our decision to focus on α-approximate pure-strategy Nash equilibria and by discussing three alternative ways of relaxing the problem.

First, we could ignore the non-existence of pure-strategy Nash equilibria and prove bounds on the price of stability for instances in which such equilibria *do* exist. This approach has been successfully applied to bounding the price of anarchy in weighted unsplittable selfish routing games [5, 9], which do not always possess pure-strategy Nash equilibria [23]. Unfortunately, for weighted Shapley network design games, a consequence of our constructions is that no *sublinear* bound on the price of stability is possible in the parameter range where pure-strategy Nash equilibria need not exist. Precisely, we show in Section 4.4 that for every function $f(x) = o(\log x/\log\log x)$, there is a family of weighted Shapley network design games in which $f(w_{max})$-approximate Nash equilibria exist, but all such equilibria have cost an $\Omega(W)$ factor times that of optimal.

Second, we could study the recent notion of "sink equilibria" due to Goemans, Mirrokni, and Vetta [13]. A sink equilibrium of a game is a strongly connected component with no outgoing arcs in the best-response graph of the game (where nodes correspond to outcomes, arcs to best-response deviations by players). Note that once a sequence of best-response deviations leads to a sink equilibrium, it will never again escape it. Sink equilibria always exist, although they can be extremely large (such as the entire best-response graph). The social value (or cost) of a sink equilibrium is defined in [13] as the expected value of a random state, where the expectation is over the stationary distribution of a random walk in the directed graph corresponding to the equilibrium. While sink equilibria are a well-motivated concept and make analyses of the price of anarchy more robust and realistic (and this was the motivation in [13]), it is not clear that they are relevant to price of stability analyses, where we envision a single solution being proposed to players as a low-cost, stable outcome. Note in particular that a sink equilibrium offers no guarantee to an individual player except for a trivial one: if a node is reached via a best-response deviation by that player, then of course it will not want to deviate again. Unfortunately, this is small consolation to a player that spends most of its time in undesirable states while other players take their turns performing their own best-response deviations.

Third, and perhaps most obviously, we could study *mixed-strategy Nash equilibria*, where each player can randomize over its path set to minimize its expected cost. Every weighted Shapley network design game admits at least one mixed-strategy Nash equilibrium by Nash's Theorem [20]. As with sink equilibria, however, it is not clear how to interpret mixed-strategy equilibria in the context of the price of stability of network design (see also the discussion in [3]). For example, a mixed-strategy Nash equilibrium could randomize only over outcomes that are not α-approximate Nash equilibria for any reasonable value of α, leading only to realizations that would be extremely difficult to enforce. One possible solution would be to implement some type of contract binding the players to the realization of a mixed-strategy Nash equilibrium. Once enforceable contracts are assumed, however, it is arguably more realistic to simply build a near-optimal network and appropriately transfer payments from players incurring small cost to those incurring large cost. Finally, if one insists on making assumptions that cause mixed-strategy Nash equilibria to be realistically implementable, then we advocate *correlated equilibria* [4] as a more suitable candidate for price of stability analyses. Correlated equilibria are no harder to justify than mixed-strategy Nash equilibria for the price of stability of network design. Moreover, since they form a convex set containing all mixed-strategy Nash equilibria, they seem likely to be both more powerful and more analytically tractable. We note that the inefficiency of correlated equilibria in different applications has largely resisted analysis so far (though see [8]), and leave this direction open for future research.

2 The Model

We now briefly formalize the model of network design with selfish players that we outlined in the Introduction. A *weighted Shapley network design game* is a directed graph $G = (V, E)$ with k source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$, where each pair (s_i, t_i) is associated with a player i that has a positive weight w_i. By scaling, we can assume that $\min_i w_i = 1$. Finally, each edge $e \in E$ has a nonnegative cost c_e.

The strategies for player i are the simple s_i-t_i paths $\mathcal{P}_i$ in G. An outcome of the game is a vector $(P_1, \ldots, P_k)$ of paths with $P_i \in \mathcal{P}_i$ for each i. For a given outcome and a player i, the cost share π_e^i of an edge $e \in P_i$ is $c_e \cdot w_i / W_e$, where $W_e = \sum_{j\,:\,e \in P_j} w_j$ is the total weight of the players that select a path containing e. The cost to player i in an outcome is the sum of its cost shares: $c_i(P_1, \ldots, P_k) = \sum_{e \in P_i} \pi_e^i$.

An outcome $(P_1, \ldots, P_k)$ is a *(pure-strategy) Nash equilibrium* if, for each i, P_i minimizes c_i over all paths in $\mathcal{P}_i$ while keeping P_j fixed for $j \neq i$. An outcome $(P_1, \ldots, P_k)$ is an *α-approximate Nash equilibrium* if for each i,

$$c_i(P_1, \ldots, P_i, \ldots, P_k) \leq \alpha \cdot c_i(P_1, \ldots, P_i', \ldots, P_k)$$

for all $P_i' \in \mathcal{P}_i$.

The *cost* $C(P_1, \ldots, P_k)$ of an outcome $(P_1, \ldots, P_k)$ is defined as $\sum_{e \in \cup_i P_i} c_e$. The *price of stability* of a game that has at least one Nash equilibrium is $C(N)/C(O)$, where N is a Nash equilibrium of minimum-possible cost and O is an outcome of minimum-possible cost. The *price of stability of α-approximate Nash equilibria* is defined analogously. Finally, we will sometime use the expression (α, β)*-approximate Nash equilibrium* to mean an outcome that is an α-approximate Nash equilibrium and that has cost at most a β factor times that of optimal.

3 Non-Existence of Equilibria with Weighted Players

In this section, we prove that weighted Shapley network design games need not possess a pure-strategy Nash equilibrium.

Proposition 3.1 *There is a 3-player weighted Shapley network design game that admits no pure-strategy Nash equilibrium. Moreover, the underlying network is undirected with a single sink, and the ratio between the maximum and minimum player weights can be made arbitrarily small.*

Recall that Anshelevich et al. [2] proved that every two-player weighted Shapley network design game has a pure-strategy Nash equilibrium.

Proof of Proposition 3.1: We first present a directed network with no pure-strategy Nash equilibrium and then describe how to convert it into an undirected example. The directed version is shown in Figure 1. Let G denote this graph and $w > 1$ a parameter. The players with sources s_1, s_2, and s_3 have weights w^2, 1, and w, respectively. All three players share a common sink t. Costs for the edges of G are defined as in Table 1, where we assume that $\epsilon > 0$ is much smaller than $1/w^3$.

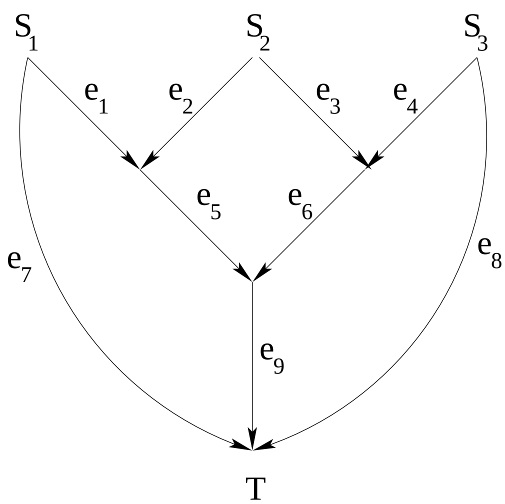

Figure 1: A three-player weighted Shapley network design game with a single-sink network and no pure-strategy Nash equilibrium.

Let c_i denote the cost of edge e_i. Our argument will rely on the following two chains of inequalities, which follow from our choice of edge costs:

$$c_5 \cdot \frac{w^2}{w^2+1} \; + \; c_9 \cdot \frac{w^2}{w^2+1} \; > \; c_7 \; > \; c_5 \; + \; c_9 \cdot \frac{w^2}{w^2+w+1}; \tag{1}$$

and

$$c_6 + \; c_9 \cdot \frac{w}{w^2+w+1} \; > \; c_8 > \; c_6 \cdot \frac{w}{w+1} \; + \; c_9 \cdot \frac{w}{w+1}. \tag{2}$$

Edge	Cost	Edge	Cost	Edge	Cost
e_1	0	e_2	3ϵ	e_3	0
e_4	0	e_5	$w^3/(w^2+w+1)-\epsilon$	e_6	$w^3/(w^2+w+1)+\epsilon$
e_7	$[(w^3+w^2)/(w^2+w+1)]$ $-[\epsilon(2w^2+1)/(2w^2+2)]$	e_8	$[(w^3+w)/(w^2+w+1)]$ $+[\epsilon(2w+1)/(2w+2)]$	e_9	1

Table 1: Edge costs for the graph G in Proposition 3.1.

(For the reader who wishes to verify these, we suggest initially taking $w = 2$.)

Now suppose for contradiction that a (pure-strategy) Nash equilibrium exists in G. Suppose further that the second player uses the path $e_2 \to e_5 \to e_9$ in this equilibrium. The first half of the inequality (2) implies that the third player must be using the one-hop path e_8 (it would share edge e_6 with no other player, and in the best case would share edge e_9 with both of the other players). The first half of inequality (1) then implies that the first player must use the one-hop path e_7. But then the second player would prefer the path $e_3 \to e_6 \to e_9$, contradicting our initial assumption.

Similarly, if the second player uses the path $e_3 \to e_6 \to e_9$ in a Nash equilibrium, then the second half of inequality (2) implies that the third player must be using the path $e_4 \to e_6 \to e_9$. The second half of inequality (1) then implies that the first player must use $e_1 \to e_5 \to e_9$. Since this would cause the path $e_2 \to e_5 \to e_9$ to be preferable to the second player, we again arrive at a contradiction. There is thus no Nash equilibrium in this weighted Shapley network design game.

To convert this directed example into an undirected one, simply make all of the edges undirected and add a large constant $M >> w^3$ to the costs of the edges e_1, e_2, e_3, e_4, e_7, and e_8. The cost of every path in the original directed network increases by exactly M; the cost of new paths are at least $2M$. As long as M is sufficiently large, no player will use one of the new undirected paths in an equilibrium, and all of the arguments for the directed network carry over without change. ■

4 Low-Cost Approximate Nash Equilibria: Lower Bounds

In this section we present negative results on the existence and price of stability of α-approximate Nash equilibria in weighted Shapley network design games. We state our lower bound on the feasible trade-offs between cost and stability in Subsection 4.1. The technical heart of this lower bound is Subsection 4.3, where we construct weighted Shapley network design games without $o(\log w_{max}/\log\log w_{max})$-approximate Nash equilibria. To illustrate our main ideas, we present a simpler version of this construction in Subsection 4.2. Finally, Section 4.4 proves that even when $o(\log w_{max}/\log\log w_{max})$-approximate Nash equilibria exist, such equilibria can have arbitrarily large cost.

We will give nearly matching positive results in Section 5.

4.1 Lower Bounds for Trading Stability for Cost

The goal of this section is to establish the following lower bound on the feasible trade-offs between the stability and the cost of approximate Nash equilibria: for every $\alpha = \Omega(\log w_{max} / \log \log w_{max})$, a price of stability of $O((\log W)/\alpha)$ can be achieved only by relaxing equilibrium constraints by an $\Omega(\alpha)$ factor. Precisely, we will prove the following.

Theorem 4.1 *Let f and g be two bivariate real-valued functions, increasing in each argument, such that every weighted Shapley network design game with maximum player weight w_{max} and sum of player weights W admits an $f(w_{max}, W)$-approximate Nash equilibrium with cost no more than a $(1 + g(w_{max}, W))$ factor times that of optimal. Then:*

(a) for some constant c,

$$f(w_{max}, W) \geq c \frac{\log w_{max}}{\log \log w_{max}}$$

for all $W \geq w_{max} \geq 1$;

(b) for some constant c,

$$f(w_{max}, W) \cdot g(w_{max}, W) \geq c \log W$$

for all $W \geq w_{max} \geq 1$.

As we will see in the next section, Theorem 4.1 is optimal up to a doubly logarithmic factor in part (a).

4.2 Networks Without Approximate Nash Equilibria

Our proof of Theorem 4.1 is fairly technical. To introduce the main ideas in the proof, we first briefly describe a simpler family of networks. These networks can be used to define weighted Shapley network design games without $(2 - \epsilon)$-approximate Nash equilibria, where $\epsilon > 0$ is arbitrarily small. Since proving this fact is not easy, we discuss only the construction. In Section 4.3 we build on this construction to prove Theorem 4.1.

We consider the network and source-sink pairs shown in Figure 2. (Each source of the form $s_{0,j}$ corresponds to the sink t_0.) In the figure, all sources and sinks have only one incident arc, except for s^* and $\bar{s}^*$, which each have one incoming and one outgoing arc. There are two *primary paths*, which contain all of the edges on the lower and upper horizontal paths, respectively. Loosely speaking, each player chooses between two "short" paths (one for each primary path), and long paths that "wrap around" the network and intersect both primary paths. For suitable choices of edge costs and player weights, long paths will not be used in any approximate Nash equilibrium. Edges not on either primary path have cost 0. We refrain from precisely specifying the costs of other edges or the players' weights; roughly, the former quantity increases exponentially while the latter quantity decreases exponentially from "left" to "right" in Figure 2.

The plan for proving that these networks do not have approximate Nash equilibria is as follows. The player with source-sink pair (s_i, t_i), which has the largest weight, must choose

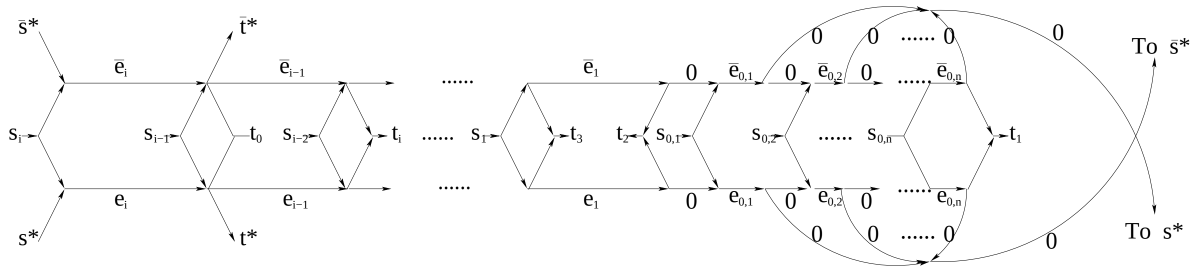

Figure 2: A network with no $(2-\epsilon)$-approximate Nash equilibria. The two primary paths are shown in bold.

one of the primary paths. This decision makes the edges on this path look cheap to the other players. Second, whichever primary path the largest player chooses, its decision must cascade through the rest of the players. Third, the n players with sink t_0 then wrap around to the other primary path, which in turn causes the largest player to want to switch to the other primary path, thereby precluding any stable outcome. We give a rigorous version of this argument for a more complex network in the next section.

4.3 Networks Without $o(\log w_{max}/\log\log w_{max})$-Approximate Nash Equilibria

We next build on the construction in the previous section to show a near-optimal lower bound on the existence of approximate Nash equilibria.

Theorem 4.2 *For every function $f(x)=o(\log x/\log\log x)$, there is a family of weighted Shapley network design games that do not admit $f(w_{max})$-approximate Nash equilibria as $w_{max}\to\infty$.*

The high-level idea behind the proof of Theorem 4.2 is similar to that of the previous construction, with upper and lower primary paths that wrap around and cross over at their ends. As before, only edges on the primary paths have nonzero cost and most players choose between short paths on the upper and lower primary paths.

The source of amplification in the new construction is that, instead of having a sequence of players with exponentially decreasing weights, we will use a group of players in each weight class. For each stage of the network, there will be a corresponding sequence of edges on each of the primary paths instead of just one. The details follow.

Set the parameter p to the square of a sufficiently large integer. Let α denote $\mathcal{H}_{\sqrt{p}}$, where $\mathcal{H}_j=\sum_{\ell=1}^{j}1/\ell\approx\ln j$ is the jth Harmonic number; this will be roughly our lower bound on the approximation factor necessary to guarantee the existence of approximate Nash equilibria. Set i to $\lceil 5\log_2\alpha\rceil+2$ and n to $2p^{2i}\alpha$. We consider a network that comprises $i-1$ stages that are connected in series. All stages but the first and last have the structure shown in Figure 3(a). The first and last stages are depicted in Figure 3(b) and (c), respectively.

Primary paths are defined as in the previous section. A path is *short* if it contains edges from only one of the primary paths, and is *long* otherwise. We further classify a short path as *upper* or *lower*, depending on which of the two primary paths it intersects.

The cost of the edges are:

- $c(e_{2i}) = c(\overline{e}_{2i}) = p^{2i}$;
- $c(e_{2i-2}) = c(\overline{e}_{2i-2}) = p^{2i}/\alpha$;
- $c(e_{2i-1}) = c(\overline{e}_{2i-1}) = 3p^{2i}\alpha^3$;
- $c(e_{2j,\ell}) = c(\overline{e}_{2j,\ell}) = 2^{i-j-1}p^{2i}/\ell\alpha^2$ for $j = 1, 2, \ldots, i-2$ and $\ell = 1, 2, \ldots, \sqrt{p}$;
- $c(e_{2j-1}) = c(\overline{e}_{2j-1}) = 2^{i-j-1}p^{2i+\frac{1}{2}}$ for $j = 2, 3, \ldots, i-1$;
- $c(e_{0,j}) = c(\overline{e}_{0,j}) = \alpha$, for $j = 1, 2, \ldots, n$;
- all other edges have cost 0.

The players in the network game are as follows. Every player will be classified as either *even*, *odd*, or *small*.

- Players A_{2i}, A^*, and $\bar{A}^*$ (with corresponding source-sink pairs (s_i, t_i), (s^*, t^*), and $(\overline{s}^*, \overline{t}^*)$) have weight p^{2i}. Player A_{2i} is an even player; the other two are odd.
- For each $j = 3, 4, \ldots, i-1$ and $\ell = 1, 2, \ldots, \sqrt{p}$, there is an even player $A_{2j,\ell}$ with weight p^{2j} and source-sink pair $(s_{2j}, t_{2j,\ell})$.
- For each $j = 1, 2, \ldots, i-1$, there are two odd players A_{2j+1} and $\overline{A}_{2j+1}$ with source-sink pairs (s_{2j+1}, t_{2j+1}) and $(\overline{s}_{2j+1}, \overline{t}_{2j+1})$, respectively, and with weight p^{2j+1}.
- There are even players A_4 and A_2 with respective weights p^4 and p^2 and respective source-sink pairs (s_4, t_4) and (s_2, t_2).
- For each $\ell = 1, 2, \ldots, n$, there is a small player $A_{0,\ell}$ with weight 1 and source-sink pair $(s_{0,\ell}, t_0)$.

We now give the proof.

Proof of Theorem 4.2: Consider the weighted Shapley network design game described above, where the parameter p is sufficiently large. We begin with a few preliminary observations. First, the maximum and minimum player weights are p^{2i} and 1, respectively. Thus $w_{max} = p^{\Theta(\log\log p)}$ while $\alpha = \mathcal{H}_{\sqrt{p}} = \Theta(\log w_{max}/\log\log w_{max})$. Second, odd players have only one available (simple) path. Third, the sum W of the player weights is

$$3p^{2i} + \sqrt{p}\sum_{j=3}^{i-1} p^{2j} + 2\sum_{j=1}^{i-1} p^{2j+1} + p^4 + p^2 + 2p^{2i}\alpha \leq 3p^{2i}\alpha \tag{3}$$

for p sufficiently large. We now establish the following six claims in turn.

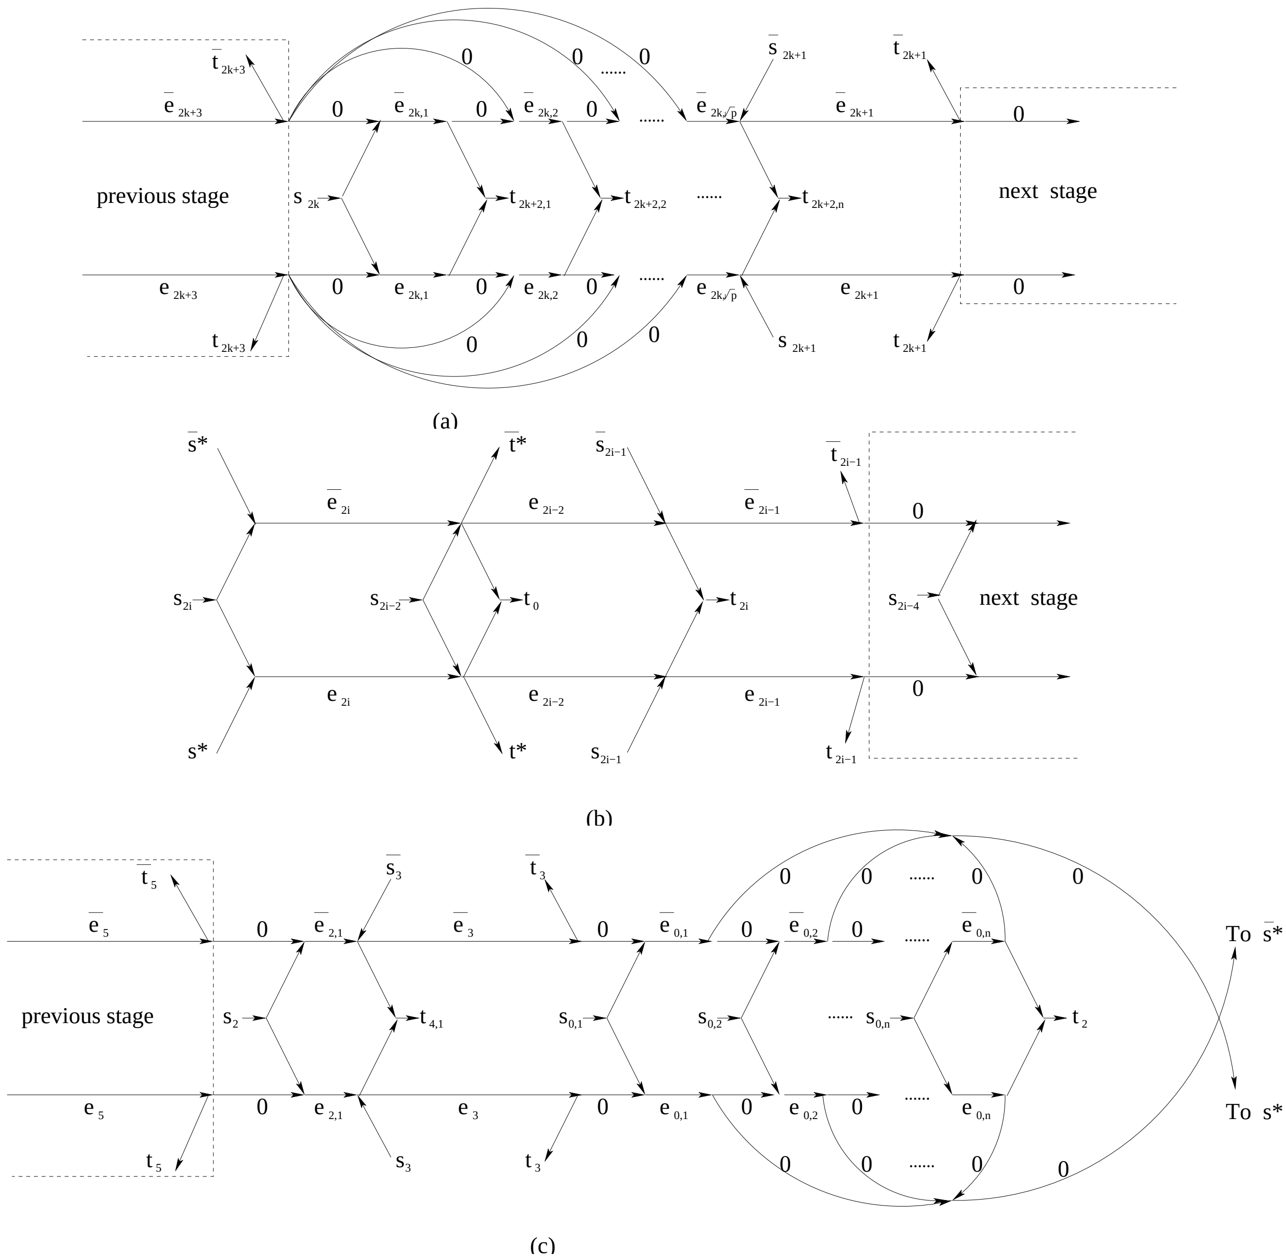

Figure 3: (a) The structure of the $(i-k)$-th stage. (b) The structure of the first stage. (c) The structure of the last stage.

(C1) In every $\alpha/6$-approximate Nash equilibrium, no small player uses a path that contains both of the edges e_{2i} and $\overline{e}_{2i}$.

(C2) In every $\alpha/6$-approximate Nash equilibrium, player A_{2i} uses a short path.

(C3) In every $\alpha/6$-approximate Nash equilibrium, every even player uses a short path.

(C4) In every $\alpha/6$-approximate Nash equilibrium in which player A_{2i} uses its lower (upper) short path, all of the even players also use lower (upper) short paths.

(C5) In every $\alpha/6$-approximate Nash equilibrium in which player A_{2i} uses its lower (upper) short path, all of the small players use paths that include edge $\overline{e}_{2i}$ (e_{2i}).

(C6) In every $\alpha/6$-approximate Nash equilibrium in which each of the small players chooses a path that contains $\overline{e}_{2i}$ (e_{2i}) but not e_{2i} ($\overline{e}_{2i}$), player A_{2i} chooses its upper (lower) short path.

Since $\alpha/6 = \Theta(\log w_{max}/\log\log w_{max})$ and claims (C1), (C5), and (C6) cannot simultaneously hold, claims (C1)–(C6) imply the theorem.

To prove (C1), note that if a small player $A_{0,\ell}$ uses a path that includes both e_{2i} and $\overline{e}_{2i}$, then it traverses either edge e_{2i-1} or edge $\overline{e}_{2i-1}$. Since the cost of each of these edges is $3p^{2i}\alpha^3$, the player incurs cost at least $3p^{2i}\alpha^3/W$, where W is the sum of the players' weights. (Recall this player has unit weight.) By (3), this is at least α^2. On the other hand, if the player $A_{0,\ell}$ chooses a path containing only the non-zero cost edges $e_{0,j}$ and $\overline{e}_{2i}$ or $\overline{e}_{0,j}$ and e_{2i}, then it incurs cost at most $\alpha + p^{2i}/(1+p^{2i}) < \alpha + 1 < 2\alpha$. (The player will share edge e_{2i} or $\overline{e}_{2i}$ with the player A^* or $\overline{A}^*$, respectively.) Thus in every $\alpha/6$-approximate Nash equilibrium, no small player uses a path containing both e_{2i} and $\overline{e}_{2i}$.

The proof of (C2) is similar. If player A_{2i} does not use a short path, then it uses a path that contains either edge e_{2i-1} or edge $\overline{e}_{2i-1}$ and incurs cost at least $3p^{2i}\alpha^3 \cdot (p^{2i}/W) \geq p^{2i}\alpha^2$. If it uses a short path, then its incurred cost is at most $p^{2i}(1+1/\alpha) < 2p^{2i}$.

For claim (C3), we first prove the assertion for all even players of the form $A_{2j,\ell}$, by downward induction on j. For the base case, consider a player $A_{2i-2,\ell}$ for arbitrary $\ell \in \{1, 2, \ldots, \sqrt{p}\}$. Player A_{2i} must use either edge e_{2i-2} or $\overline{e}_{2i-2}$. The odd players A_{2i-1} and $\overline{A}_{2i-1}$ must occupy the edges e_{2i-1} and $\overline{e}_{2i-1}$. Thus, there is a short path available to player $A_{2i-2,\ell}$ with cost at most

$$\frac{p^{2i}}{\alpha}\left(\frac{p^{2i-2}}{p^{2i-2}+p^{2i}}\right) + 3p^{2i}\alpha^3\left(\frac{p^{2i-2}}{p^{2i-2}+p^{2i-1}}\right) + \sum_{m=1}^{\ell}\frac{2p^{2i}}{m\alpha^2} \leq \frac{p^{2i-2}}{\alpha} + 3p^{2i-1}\alpha^3 + \frac{2p^{2i}}{\alpha} \leq \frac{3p^{2i}}{\alpha}$$

for p sufficiently large. On the other hand, every long path of player $A_{2i-2,\ell}$ contains either edge e_{2i-3} or edge $\overline{e}_{2i-3}$. By (C1) and (C2), the total weight on this edge is at most the weight of the corresponding odd player plus the total weight of the even players other than A_{2i}, which is at most

$$p^{2i-3} + \sqrt{p}\sum_{j=3}^{i-1} p^{2j} + p^4 + p^2 \leq 2p^{2i-(3/2)} \tag{4}$$

for p sufficiently large. Therefore, if player $A_{2i-2,\ell}$ chooses a long path, it incurs cost at least

$$p^{2i+(1/2)}\left(\frac{p^{2i-2}}{p^{2i-2}+2p^{2i-(3/2)}}\right) > \frac{p^{2i}}{2}. \tag{5}$$

Inequalities (4) and (5) imply claim (C3) for players of the form $A_{2i-2,\ell}$.

For the inductive step, fix $j \in \{3, 4, \ldots, i-2\}$ and assume that every player of the form $A_{2j',\ell}$ with $j' > j$ uses a short path. Consider a player $A_{2j,\ell}$ for some $\ell \in \{1, 2, \ldots, \sqrt{p}\}$. Arguing as in the base case, the player can choose a short path and incur cost at most

$$\sum_{m=1}^{\ell} \frac{2^{i-j-1}p^{2i}}{m\alpha^2} + p^{2i+(1/2)}\frac{p^{2i-4}}{p^{2i-4}+p^{2i-3}} + \frac{2^{i-j}p^{2i}}{\ell\alpha^2} \leq \frac{2^{i-j-1}p^{2i}}{\alpha} + p^{2i-(1/2)} + \frac{2^{i-j}p^{2i}}{\alpha^2} < \frac{2^{i-j}p^{2i}}{\alpha},$$

provided p is sufficiently large. On the other hand, every long path contains either edge e_{2j-1} or edge $\overline{e}_{2j-1}$. By (C1), (C2), and the inductive hypothesis, the total weight on this edge is at most

$$p^{2j-1} + \sqrt{p}\sum_{m=3}^{j} p^{2m} + p^4 + p^2 \leq 2p^{2j+(1/2)}$$

for p sufficiently large. Thus, if player $A_{2j,\ell}$ chooses a long path, it incurs cost at least

$$2^{i-j-1}p^{2i+(1/2)}\left(\frac{p^{2j}}{p^{2j}+2p^{2j+(1/2)}}\right) > 2^{i-j-2}p^{2i}.$$

This completes the inductive step.

Finally, we establish (C3) for players A_4 and A_2. Given that (C3) holds for all other players, player A_4 has a short path on which it would incur cost at most

$$\frac{2^{i-3}p^{2i}}{\alpha} + 2^{i-2}p^{2i-(1/2)} + \frac{2^{i-2}p^{2i}}{\alpha^2} < \frac{2^{i-2}p^{2i}}{\alpha},$$

while every long path contains either e_3 or $\overline{e}_3$ and causes the player to incur cost at least

$$2^{i-2}p^{2i+(1/2)}\left(\frac{p^4}{p^4+p^3+p^2}\right) > 2^{i-3}p^{2i+(1/2)}$$

for large p. For player A_2, previous steps imply that one of the edges $e_{2,1}, \overline{e}_{2,1}$ contains player A_4 while the other is unoccupied by other players. Long paths contain both of these, so if player A_2 chooses one of them it incurs cost at least $2^{i-2}p^{2i+(1/2)}$. On the other hand, there is a short path with cost at most

$$\frac{2^{i-2}p^{2i}}{\alpha^2} + 2^{i-2}p^{2i-(3/2)} + 2p^{2i}\alpha^2 < 2^{i-2}p^{2i}\alpha^2.$$

Assuming p is large, this implies that player A_2 must choose a short path, completing the proof of (C3).

For (C4), by symmetry we can assume that player A_{2i} takes its lower short path. We proceed by contradiction. Among all even players that choose an upper short path, select the "leftmost" one—the one with maximum index j and, subject to this, with maximum index ℓ. Suppose this player is $A_{2j,\ell}$ with $j \geq 3$. If $j = i-1$, then the cost to this player on its minimum-cost lower short path is at most

$$\frac{p^{2i-2}}{\alpha} + 3p^{2i-1}\alpha^3 + \frac{2p^{2i}}{\ell\alpha^2} < \frac{4p^{2i}}{\ell\alpha^2}$$

for p large. The key point is this: by (C3) and the definition of ℓ, the only players eligible for using edge $\overline{e}_{2i-2}$ are $A_{2j,1}, \ldots, A_{2j,\ell}$. Thus, if $A_{2j,\ell}$ chooses an upper short path, it incurs cost at least $p^{2i}/\ell\alpha$, providing a contradiction.

Similarly, if the player is $A_{2j,\ell}$ with $j \in \{3, 4, \ldots, i-2\}$, then our choice of j ensures that the cost to the player on its minimum-cost lower short path is at most

$$\frac{2^{i-j-1}p^{2i-2}}{\alpha} + p^{2i-(1/2)} + \frac{2^{i-j}p^{2i}}{\ell\alpha^2} < \frac{2^{i-j+1}p^{2i}}{\ell\alpha^2},$$

while our choice of ℓ ensures that the cost incurred on every upper short path is at least $2^{i-j-1}p^{2i}/\ell\alpha$, a contradiction. Finally, suppose that the "leftmost" even player choosing a short upper path is A_4 or A_2. The contradiction in the former case is essentially the same as that for the previous case of a player $A_{2j,\ell}$ with $j \in \{3, 4, \ldots, i-2\}$. In the latter case, the cost incurred by player A_2 on its lower short path is at most

$$\frac{2^{i-2}p^{2i-2}}{\alpha^2} + 2^{i-2}p^{2i-(1/2)} + 2p^{2i}\alpha^2 < 4p^{2i}\alpha^2$$

for p large. Since edge $\overline{e}_{2,1}$ is occupied by no other player, our choice of $i = \lceil 5\log_2 \alpha \rceil + 2$ ensures that the cost incurred by player A_2 on its upper short path is at least

$$\frac{2^{i-2}p^{2i}}{\alpha^2} \geq p^{2i}\alpha^3,$$

which completes the proof of (C4).

We prove claim (C5) by contradiction. Assume that player A_{2i} chooses its lower short path. Let ℓ be the minimum index for which player $A_{0,\ell}$ chooses a path containing the edge e_{2i}. By (C1)–(C4) and our choice of ℓ, this player incurs the full α cost of edge $\overline{e}_{0,\ell}$. On the other hand, if the player chooses the $s_{0,\ell}$-t path containing $e_{0,\ell}$ and $\overline{e}_{2i}$ (and no other edges with non-zero cost), then it shares the former edge with player A_2 and incurs cost at most

$$\frac{\alpha}{1+p^2} + \frac{p^{2i}}{1+p^{2i}} < 2.$$

This completes the contradiction and the proof of (C5).

Finally, assume the hypothesis in claim (C6) holds. By (C1)–(C3), if player A_{2i} chooses its lower short path, then it shares edge e_{2i} only with player A^* and incurs cost at least $p^{2i}/2$. On the other hand, the cost of its upper short path is

$$p^{2i}\left(\frac{p^{2i}}{2p^{2i} + 2p^{2i}\alpha}\right) + \frac{p^{2i}}{\alpha}\left(\frac{p^{2i}}{p^{2i} + p^{2i-1}}\right) < \frac{2p^{2i}}{\alpha},$$

which completes the proof of (C6) and of the theorem. ■

With Theorem 4.2 in hand, we can easily finish the proof of Theorem 4.1.

Proof of Theorem 4.1: Part (a) follows immediately from Theorem 4.2. Part (b) holds even for the special case of unweighted Shapley network design games and follows from a minor modification of an example in [2]. Specifically, Anshelevich et al. [2] presented an unweighted Shapley network design game in which the minimum-cost solution has cost $1+\epsilon$, where $\epsilon > 0$ is arbitrarily small, and the unique Nash equilibrium has cost $\mathcal{H}_k$. Moreover, these two outcomes use disjoint edge sets. For each fixed value of W, we can take this example with $k = \lfloor W \rfloor$ players and scale down the costs of the edges used by the Nash equilibrium by an $f(1, W)$ factor. This yields an (unweighted) game in which the only $f(1, W)$-approximate Nash equilibrium has cost $\Omega(\log W/f(1, W))$; there is still a solution with cost $1+\epsilon$. Thus $f(1, W) \cdot g(1, W) = \Omega(\log W)$ for all $W \geq 1$. ■

4.4 A Lower Bound on the Price of Stability

In this section, we employ the networks of Section 4.3 to show that, in addition to the evaporation of α-approximate Nash equilibria once $\alpha = o(\log w_{max}/\log\log w_{max})$, in instances where such equilibria *do* exist, their cost can be extremely high. This stands in contrast to recent work on routing games [5], where there are good upper bounds on the price of anarchy even in classes of networks where Nash equilibria are not guaranteed to exist.

Proposition 4.3 *For every function $f(x) = o(\log x/\log\log x)$, there are weighted Shapley network design games that admit pure-strategy Nash equilibria, but in which all $f(w_{max})$-approximate Nash equilibria have cost $\Omega(W)$ times that of optimal.*

Proof: We adopt the notation used in Section 4.3. For every function $f(x) = o(\log x/\log\log x)$, we can find a sufficiently large constant p such that the corresponding weighted Shapley network design game G constructed in that section has no $f(w_{max})$-approximate Nash equilibria. We note that the sum of the edge costs in G is at most p^{2i+1}, provided p is sufficiently large.

We then construct a new network game as follows. We take m copies $G_1, G_2, \ldots, G_m$ of G and remove the player A^* in each of them. As shown in Figure 4, we add the extra nodes $N_1, N_2, \ldots, N_m$, N_T and T. We add the following edges:

- one edge from N_T to T with cost $C = mp^{4i}$;
- for each $j = 1, 2, \ldots, m$, one edge from N_j to T with cost C;
- for each $j = 1, 2, \ldots, m$, one zero-cost edge from N_j to G_j's copy of the vertex s^*;
- for each $j = 1, 2, \ldots, m$, one zero-cost edge from G_j's copy of the vertex t^* to N_T.

We also add m players $B_1, B_2, \ldots, B_m$. For each j, the player B_j has weight p^{2i} and source-sink pair (N_j, T). By construction, the only players that can use edges inside the network G_j are the players internal to this game and the new player B_j.

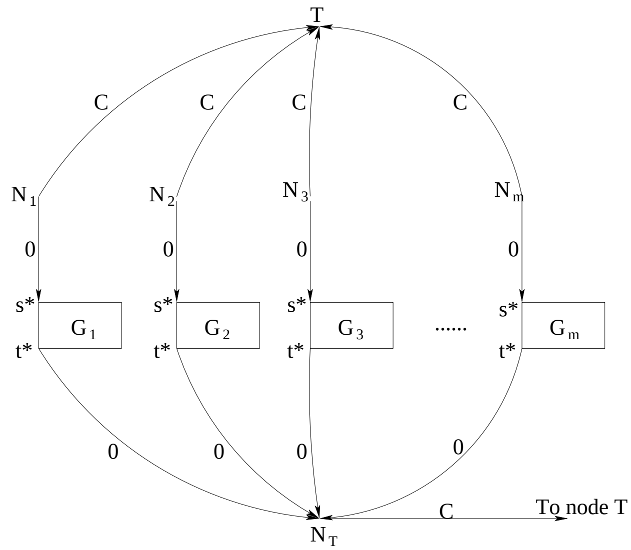

Figure 4: A network in which all approximate Nash equilibria have cost far from optimal. Every rectangle represents a network of the type described in Section 4.3.

First suppose that a player B_j connects to T via the vertices s^* and t^* internal to G_j. By the argument in the proof of Theorem 4.2, in every such outcome, some player internal to the network G_j has a deviation that decreases its cost by more than an $f(w_{max})$ factor; here, B_j is playing the role of the deleted player A^* in the game G_j. Thus no such outcome is an $f(w_{max})$-approximate Nash equilibrium of the game. On the other hand, one such outcome—with each B_j choosing the path that intersects the network G_j—has cost at most $C + mp^{2i+1}$, provided p is sufficiently large.

Now suppose that each player B_j avoids G_j and connects directly to T. No player of the form B_j can profitably deviate, since it would bear the full cost of the edge from N_T to T. Moreover, consider the following strategies for the players internal to a game G_j. Each even player of G_j chooses its minimum-cost upper short path. Each small player of G_j chooses its minimum-cost path that includes the edge $\overline{e}_{2i}$—beginning on the lower primary path and wrapping around to the upper one. We claim that in this case, no player internal to G_j has an incentive to deviate. This claim is easy to see for the even players. (Recall odd players only have one available strategy.) No small player wants to deviate; since player B_j is avoiding the game G_j, the edge e_{2i} is unoccupied, and every deviation that includes it would cost a small player the full p^{2i} amount.

In conclusion, the network game admits a pure-strategy Nash equilibrium, but every $f(w_{max})$-approximate Nash equilibrium has cost at least mC, which is an $\Omega(m)$ factor times larger than the optimal cost. Since m can be arbitrarily large, the proposition follows. ■

5 Low-Cost Approximate Equilibria: Upper Bounds

In this section we prove our main positive result, that every weighted Shapley network design game admits an approximate Nash equilibrium with low cost. Specifically, we show that for all $\alpha = \Omega(\log w_{max})$, every such game admits an $O(\alpha)$-approximate Nash equilibrium with cost an $O((\log W)/\alpha)$ times that of optimal. (Recall that w_{max} and W denote the maximum player weight and the sum of the players' weights, respectively.) In particular, every weighted Shapley network design game possesses an $O(\log W)$-approximate Nash equilibrium with cost at most a constant times that of optimal. This is a new result even for unweighted Shapley network design games.

At a high level, our proof is based on the "potential function method" that has been previously used to bound the price of anarchy and stability in a number of different games (see [24]). A real-valued function Φ defined on the outcomes of a game is a *potential function* if, for every player i and every possible deviation by that player, the change in the value of Φ equals the change in player i's objective function. Thus a potential function "tracks" successive deviations by players. In particular, local optima of a potential function are precisely the pure-strategy Nash equilibria of the game. Potential functions were originally applied in noncooperative game theory by Beckmann, McGuire, and Winsten [7], Rosenthal [22], and Monderer and Shapley [17], in successively more general settings, to prove the existence of Nash equilibria. Potential functions can also be used to bound the price of stability: if a game has a potential function Φ that is always close to the true social cost, then a global optimum of Φ, or any local optimum reachable from the min-cost outcome via best-response deviations, has cost close to optimal. Indeed, Anshelevich et al. [2] proved both the existence of Nash equilibria and an $\mathcal{H}_k$ upper bound on the price of stability in unweighted Shapley network design games using a potential function.

Proposition 3.1 implies that weighted Shapley network design games do not generally admit a potential function (see also [2]). We nonetheless show that ideas from potential functions can be used to derive a nearly optimal stability vs. cost trade-off for approximate Nash equilibria of weighted Shapley network design games. The initial idea is simple: we identify an "approximate potential function", which decreases whenever a player deviates and decreases its cost by a sufficiently large factor. This argument will imply the existence of an $O(\log w_{max})$-approximate Nash equilibrium with cost within an $O(\log W)$ factor of optimal in every weighted Shapley network design game.

Extending this argument to obtain a stability vs. cost trade-off requires further work. The reason is that we will use a common approximate potential function for all points on the trade-off curve, and this potential function can overestimate the true cost by as much as a $\Theta(\log W)$ factor. This function therefore seems incapable of proving an $o(\log W)$ approximation factor for the cost, even if we relax equilibrium constraints by a large factor. We overcome this problem by more carefully considering how extra cost is incurred throughout best-response dynamics starting from a minimum-cost outcome. More precisely, we show that as we increase the relaxation factor on the equilibrium constraints, the allowable best-response deviations lead to more rapid decreases in the value of our approximate potential function. The formal statement is as follows (cf., Theorem 4.1).

Theorem 5.1 *Let f and g be two bivariate real-valued functions satisfying:*

(a)

$$f(w_{max}, W) \geq 2\log_2[e(1 + w_{max})] \tag{6}$$

for all $W \geq w_{max} \geq 1$; and

(b)

$$f(w_{max}, W) \cdot g(w_{max}, W) \geq 2\log_2(1 + W)$$

for all $W \geq w_{max} \geq 1$.

Then every weighted Shapley network design game with maximum player weight w_{max} and sum of player weights W admits an $f(w_{max}, W)$-approximate Nash equilibrium with cost at most $(1 + g(w_{max}, W))$ times that of optimal.

Before proving the theorem, we establish some preliminary results.

Fact 5.2 *Let x and y be real numbers, and suppose that $y \geq 1$ and that $x = 0$ or $x \geq 1$. Then:*

(a) $\log_2(1 + x + y) - \log_2(1 + x) \geq \frac{y}{x+y}$*; and*

(b) $\log_2(1 + x + y) - \log_2(1 + x) < \log_2[e(1 + y)] \cdot \frac{y}{x+y}$.

Proof: For both parts, we will use the fact that $(1 + \frac{1}{x})^x$ approaches e monotonically from below as $x \to \infty$. For part (a), first note that if $x \geq 0$ and $y \geq 1 + x$, then the inequality holds: the right-hand side is at most 1 while the left-hand side equals $\log_2(1 + \frac{y}{1+x}) \geq 1$. So suppose that $y < 1 + x$; then

$$\left(1 + \frac{y}{1+x}\right)^{\frac{x+y}{y}} \geq \left(1 + \frac{y}{1+x}\right)^{\frac{1+x}{y}} \geq 2.$$

Raising both sides of this inequality to the $y/(x+y)$ power and then taking logarithms (base 2) verifies the claim.

For part (b), we have

$$\begin{aligned}\left(1 + \frac{y}{1+x}\right)^{\frac{x+y}{y}} &= \left(1 + \frac{y}{1+x}\right)^{\frac{1+x}{y}} \left(1 + \frac{y}{1+x}\right)^{\frac{y-1}{y}} \\ &< \left(1 + \frac{y}{1+x}\right)^{\frac{1+x}{y}} \left(1 + \frac{y}{1+x}\right) \\ &\leq e(1 + y).\end{aligned}$$

As before, raising both sides of this inequality to the $y/(x + y)$ power and then taking logarithms (base 2) verifies the claimed inequality. ■

We next consider the existence of approximate Nash equilibria without worrying about their cost.

Lemma 5.3 *For every function $f(w_{max}, W)$ satisfying*

$$f(w_{max}, W) \geq \log_2[e(1 + w_{max})] \tag{7}$$

for all $W \geq w_{max} \geq 1$, every weighted Shapley network design game admits an $f(w_{max}, W)$-approximate Nash equilibrium.

Proof: We define an approximate potential function Φ for a weighted Shapley network design game as follows: for an outcome $(P_1, \ldots, P_k)$ of the game, define

$$\Phi(P_1, \ldots, P_k) = \sum_{e \in E} c_e \log_2(1 + W_e),$$

where $W_e = \sum_{j\,:\,e \in P_j} w_j$. Call a deviation by a player from one outcome to another α-*improving* if the deviation decreases the cost incurred by the player by at least an α multiplicative factor. Thus α-approximate Nash equilibria are those outcomes from which no α-improving deviations exist. Since there are a finite number of outcomes, we can prove the lemma by showing that $f(w_{max}, W)$-improving deviations strictly decrease the approximate potential function Φ.

Consider an α-improving deviation of player i from the outcome $(P_1, \ldots, P_k)$, say to the path Q_i, where α equals $f(w_{max}, W)$. We assume that P_i and Q_i are disjoint; if this is not the case, the following argument can be applied to $P_i \setminus Q_i$ and $Q_i \setminus P_i$ instead. By the definition of α-improving, we have

$$\sum_{e \in Q_i} c_e \cdot \frac{w_i}{W_e + w_i} \leq \frac{1}{f(w_{max}, W)} \sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e}, \tag{8}$$

where $W_e = \sum_{j\,:\,e \in P_j} w_j$ denotes the total weight on edge e before player i's deviation.

We can then derive the following:

$$\begin{aligned}
\Delta\Phi &= \sum_{e \in Q_i} c_e \cdot [\log_2(1 + W_e + w_i) - \log_2(1 + W_e)] - \\
&\qquad \sum_{e \in P_i} c_e \cdot [\log_2(1 + W_e) - \log_2(1 + W_e - w_i)] \qquad (9)\\
&< \sum_{e \in Q_i} c_e \cdot \left[\log_2[e(1 + w_i)] \frac{w_i}{W_e + w_i}\right] - \sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e} \qquad (10)\\
&\leq \log_2[e(1 + w_{max})] \sum_{e \in Q_i} c_e \cdot \frac{w_i}{W_e + w_i} - \sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e}\\
&\leq -\sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e} \cdot \frac{f(w_{max}, W) - \log_2[e(1 + w_{max})]}{f(w_{max}, W)} \qquad (11)\\
&\leq 0.
\end{aligned}$$

In this derivation, the equality (9) follows from the definition of Φ; the inequality (10) follows from Fact 5.2, with Fact 5.2(b) applied to each term in the first sum with $x = W_e$

and $y = w_i$, and Fact 5.2(a) applied to each term in the second sum with $x = W_e - w_i$ and $y = w_i$; inequality (11) follows from (8); and the final inequality follows from assumption (7). ■

We now extend the argument in the proof of Lemma 5.3 to account for the cost of approximate equilibria.

Proof of Theorem 5.1: Consider a maximal sequence of $f(w_{max}, W)$-improving deviations that begins in a minimum-cost outcome with cost C^*. By Lemma 5.3, this sequence is finite and terminates at a $f(w_{max}, W)$-approximate Nash equilibrium. Consider a deviation in this sequence by a player i from a path P_i to a path Q_i, and let A denote the cost of the edges of Q_i that were previously vacant (i.e., used by no player). We then have

$$\Delta\Phi \leq -\sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e} \cdot \frac{f(w_{max}, W) - \log_2[e(1 + w_{max})]}{f(w_{max}, W)} \tag{12}$$

$$\leq -\frac{1}{2} \sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e} \tag{13}$$

$$\leq -\frac{1}{2} A \cdot f(w_{max}, W), \tag{14}$$

where inequality (12) is the same as inequality (11) in the proof of Lemma 5.3; inequality (13) follows from assumption (6); and inequality (14) follows from the fact that the cost incurred by player i before its deviation is at least $f(w_{max}, W)$ times the cost it incurs after the deviation, which is at least the sum A of the costs of the previously vacant edges.

Hence, in the maximal sequence of $f(w_{max}, W)$-improving deviations, whenever the cost of the current outcome increases by an additive factor of A, the potential function Φ decreases by at least $A \cdot f(w_{max}, W)/2$. By the definition of Φ, the potential function value of the optimal outcome is at most a $\log_2(1 + W)$ multiplicative factor larger than its cost C^*. Moreover, the potential function is always nonnegative and only decreases throughout the sequence of deviations. Therefore, the cost only increases by a $2C^* \log_2(1 + W)/f(w_{max}, W)$ additive factor throughout the entire sequence of deviations. The sequence thus terminates in a $(f(w_{max}, W), 1 + (2\log_2(1 + W)/f(w_{max}, W)))$-approximate Nash equilibrium. ■

Remark 5.4 Our proof of Theorem 5.1 is quite flexible and carries over to extensions known for the unweighted case [2]. For example, Theorem 5.1 and its proof hold for congestion games (where the strategy set of a player is an arbitrary collection of subsets of a ground set) and for concave (instead of constant) edge costs.

6 Future Directions

The present paper gives an essentially tight analysis of the feasible trade-offs between the stability and cost of approximate Nash equilibria in Shapley network design games. On the other hand, the corresponding trade-off curve in several natural special cases is not well understood. For example, what are the feasible trade-offs in undirected networks? In single-sink networks? Or when there is only a small number of distinct player weights?

Acknowledgments

We thank the anonymous referees for their comments.

References

[1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 89–98, 2006.

[2] 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 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

[3] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 511–520, 2003.

[4] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[5] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

[6] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.

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

[8] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual Eurpoean Symposium on Algorithms (ESA)*, pages 59–70, 2005.

[9] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

[10] J. Corbo and D. C. Parkes. The price of selfish behavior in bilateral network formation games. In *Proceedings of the 24th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 99–107, 2005.

[11] P. Dubey. Inefficiency of Nash equilibria. *Mathematics of Operations Research*, 11(1):1–8, 1986.

[12] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. J. Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[13] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[14] M. O. Jackson. A survey of models of network formation: Stability and efficiency. In G. Demange and M. Wooders, editors, *Group Formation in Economics; Networks, Clubs, and Coalitions*, chapter 1. Cambridge University Press, 2005.

[15] K. Jain and V. V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–372, 2001.

[16] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[17] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[18] T. Moscibroda, S. Schmid, and R. Wattenhofer. On the topologies formed by selfish peers. In *Proceedings of the 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 133–142, 2006.

[19] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[20] J. F. Nash. Equilibrium points in N-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

[21] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

[22] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[23] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.

[24] T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proceedings of the International Congress of Mathematicians (ICM)*, volume III, pages 1071–1094, 2006.

[25] A. S. Schulz and N. Stier Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 86–87, 2003.