Network Design with Weighted Players (SPAA 2006 Full Paper Submission)

Ho-Lin Chen* Tim Roughgarden†

March 7, 2006

Abstract

We consider a model of game-theoretic network design initially studied by Anshelevich et al. [1], 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. [1] proved that pure-strategy Nash equilibria always exist and that the price of stability—the ratio in costs of a minimum-cost Nash equilibrium and 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 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

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. [5, 12, 13] 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 [9].

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 [14], 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 [22] and was so-called in Anshelevich et al. [1]. The price of stability has primarily been studied in network design games [1, 2], 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 incessantly prevent the network users to 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 can be regarded as the necessary degradation in the solution quality caused by imposing the game-theoretic constraint of stability.

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 [18] and initially explored independently by Anshelevich et al. [2] and Fabrikant et al. [10]. 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. Closest to the present work is a variation on the model of [2] that was proposed and studied by Anshelevich et at. [1], which they called *network design with Shapley cost sharing* and we will abbreviate to *Shapley network design games*.

The most basic model considered in [1] 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 assumed to be 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.

Of all the ways to share the social cost among the players, 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. [16]); and, as shown in [1], it coaxes benign behavior from the players. Specifically, Anshelevich et al. [1] showed that a pure-strategy Nash equilibrium always exists—unlike with the more general cost-sharing that was allowed in the predecessor model [2]—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. [1] also provided an example showing that this upper bound is the best possible, and proved numerous extensions.

Shapley Cost Sharing with Weighted Players

A natural and important extension that Anshelevich et al. [1] 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 could arise for many reasons. For example, each player could represent the traffic of a large population, such as the customers of an Internet Service Provider, and all such populations 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. But 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. [1] 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 [1] for weighted games concerned only the special cases of 2-player games and of single-commodity games (where all players have both the same source and the same sink). 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.

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 [1].) Thus there are no large classes of weighted Shapley network design games that always possess pure-strategy Nash equilibria beyond those identified in [1].

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. [2]—where for a different but related network design game, pure-strategy Nash equilibria did not necessarily exist—and consider *approximate* pure-strategy Nash equilibria. We say that 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 more generally understand 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?

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 [1].)

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.

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 recently been successively applied to bounding the price of anarchy in weighted unsplittable selfish routing games [4, 8], which do not always possess pure-strategy Nash equilibria [11, 20]. 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 will show in the full version of the paper 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 [11]. 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 [11] 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 [11]), 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 [17]. 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 [2]). For example, a mixed-strategy Nash equilibrium could potentially 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* [3] 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 [7]), 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^i_e 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: $c_i(P_1, \ldots, P_k) = \sum_{e \in P_i} \pi^i_e$.

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* of an outcome $(P_1, \ldots, P_k)$, denoted by $C(P_1, \ldots, P_k)$, is $\sum_{e \in \cup_i P_i} c_e$. 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 Nash 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. [1] 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 $1/w^3$.

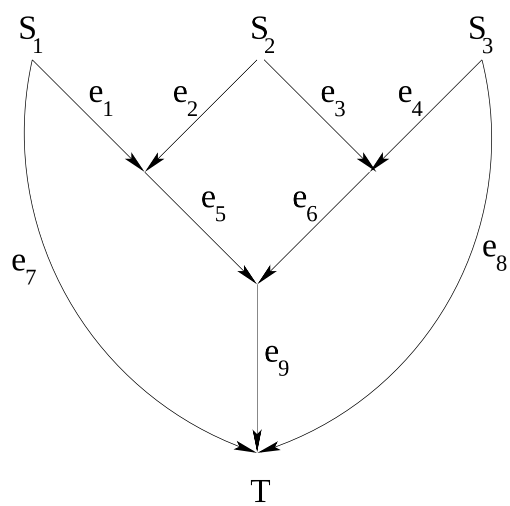

Figure 1: A three-player weighted Shapley network design game with a single-sink network and no pure-strategy Nash equilibrium.

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:

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

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.

and

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

(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 than 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. We illustrate a simpler version of this construction in Subsection 4.2, which is enough to rule out the existence of $(2-\epsilon)$-approximate Nash equilibria for arbitrarily small $\epsilon > 0$.

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 at most $(1 + g(w_{max}, W))$ 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 $(2-\epsilon)$-Approximate Nash Equilibria

We now work toward Theorem 4.1 by giving networks without α-approximate Nash equilibria for α arbitrarily close to 2. We first describe the network, then give the intuition behind the construction, and then give the details.

We will consider the network shown in Figure 2. In the figure, all sources and sinks have only one incident arc, except for s^* and $\bar{s}^*$, which each has one incoming and one outgoing arc. There are two *primary paths*, denoted Q and $\overline{Q}$, which contain all of the edges on the lower and upper horizontal paths, respectively.

There will be several parameters. We assume that we are given an arbitrarily small positive number $\epsilon_0 \leq 0.1$, with the goal of exhibiting a network with no $(2-\epsilon_0)$-approximate Nash equilibria. We then set $p = 8/\epsilon_0$, $\epsilon = \epsilon_0/8$, $i = \lceil \log_{1+\epsilon}(32p) \rceil + 2$, and $n = p^{i+1}$.

We next discuss the edge costs. To minimize subscripts, in this subsection and the next we will use $c(e)$ to denote the cost of an edge e. Edges not on either primary path have cost 0. The edge costs on the primary paths are as follows:

- $c(e_i) = c(\bar{e}_i) = \frac{p^i}{2}$;
- $c(e_j) = c(\bar{e}_j) = \frac{p^i}{8}(1+\epsilon)^{i-j-1}$, for all $j = 1, 2, \ldots, i-1$;
- $c(e_{0,j}) = c(\bar{e}_{0,j}) = 1$, for all $j = 1, 2, \ldots, n$.

The remaining edges on the primary paths have cost 0, as in Figure 2. The players are as follows.

- Players A_i, A^*, and $\bar{A}^*$ (with corresponding source-sink pairs (s_i, t_i), (s^*, t^*), and $(\bar{s}^*, \bar{t}^*)$) have weight p^i.
- For each $j = 1, 2, \ldots, i-1$, there is a player A_j with weight p^j and source-sink pair (s_j, t_j).
- There are n *small players* $A_{0,1}, A_{0,2}, \ldots, A_{0,n}$ with weight 1. For every j, the small player $A_{0,j}$ has source $s_{0,j}$ and sink t_0.

We can then prove the following.

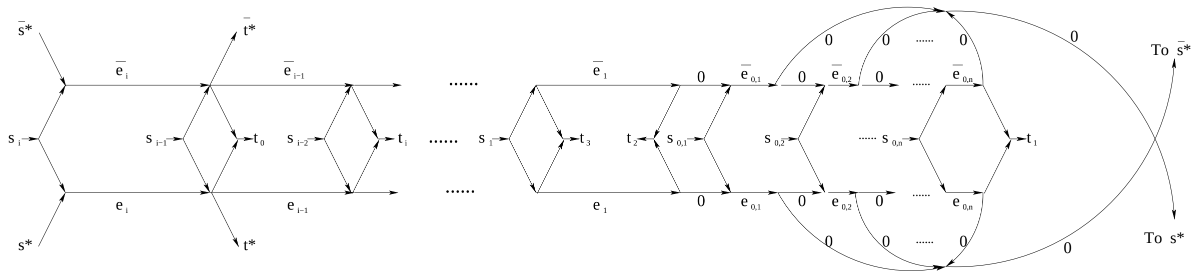

Figure 2: A network with no $(2-\epsilon)$-approximate Nash equilibria.

Theorem 4.2 *For every $\epsilon_0 \in (0, 0.1)$, the weighted Shapley network design game above has no $(2-\epsilon_0)$-approximate Nash equilibrium.*

In the proof of Theorem 4.2 we will formalize the following ideas. First, player A_i must choose one of the primary paths, which in turn makes the edges on this path look cheap to the other players. Second, whichever primary path A_i chooses, its decision must cascade through the rest of the players. Third, the n small players then wrap around to the other primary path, which in turn causes player A_i to want to switch to the other primary path, thereby precluding any stable outcome. We now make this proof approach rigorous.

Proof (sketch): We start with some terminology. A *short path* of a player that is not small is a path that leaves the player's source with a hop to one of the primary paths, follows that primary path, and then ends with a final hop to the player's destination. A *long path* of a player that is not small is one that contains edges of both primary paths. Note that for every player that is not small, all of its paths are either short or long; all such players have precisely two short paths, except for players A^* and $\bar{A}^*$, who each have one. For the proof, we will formalize the following statements in turn.

(1) In every $(2-\epsilon_0)$-approximate Nash equilibrium, no small player uses a path containing e_j or $\bar{e}_j$ with $j \in \{1, 2, \ldots, i-1\}$.

(2) In every $(2-\epsilon_0)$-approximate Nash equilibrium, player A_i uses a short path.

(3) In every $(2-\epsilon_0)$-approximate Nash equilibrium in which player A_i uses its lower (upper) short path, the players $A_1, \ldots, A_{i-1}$ also use their lower (upper) short paths.

(4) In every $(2-\epsilon_0)$-approximate Nash equilibrium in which player A_i uses its lower (upper) short path, all of the small players use paths that include the edge $\bar{e}_i$ (e_i).

(5) In every $(2-\epsilon_0)$-approximate Nash equilibrium in which all of the small players use paths that include the edge $\bar{e}_i$ (e_i), player A_i uses its upper (lower) short path.

Since (4) and (5) are mutually exclusive, proving (1)–(5) completes the proof of the theorem. Due to space constraints, we defer further details to Appendix A.1. ■

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

We next build on the construction in Theorem 4.2 to show a much stronger (and near-optimal) lower bound on the existence of approximate Nash equilibria.

Theorem 4.3 *For every function $f(x) = o(\frac{\log x}{\log\log x})$, there is a family of weighted Shapley network design games that do not admit $f(w_{max})$-approximate Nash equilibria as $w_{max} \to \infty$.*

Due to space constraints, we only describe some of the intuition behind Theorem 4.3 and defer all of the details to the Appendix. The high-level idea is similar to the previous construction, with an upper and lower primary path that wrap around and cross over at their ends. As before, only edges on the primary paths have nonzero cost and most players can choose between short paths on the upper and lower primary paths.

The source of amplification in the new construction is that, instead of having one player with weight p^i for each i as in the previous example, we will use $\sqrt{p}$ players with weight p^{2i} for each i. For each stage, there will be $\sqrt{p}$ edges on each of the main paths instead of just one.

The structure of the argument that there is no α-approximate Nash equilibrium then consists of verifying the analogous statements (1)–(5). While the proof is complicated by the larger value of α and the increased number of players and paths, it is conceptually very similar to the proof of Theorem 4.2. The details can be found in the Appendix.

With Theorem 4.3 in hand, we can easily finish the proof of Theorem 4.1.

Proof of Theorem 4.1: Part (a) follows immediately from Theorem 4.3. Part (b) holds even for the special case of unweighted Shapley network design games and follows from a minor modification of an example in [1]. Specifically, Anshelevich et al. [1] presented an unweighted Shapley network design game in which a minimum-cost solution has cost 1 and the unique Nash equilibrium has cost $\mathcal{H}_k$. Moreover, the two outcomes use disjoint edge sets. For each fixed value of W, we take this example with $k = \lfloor W \rfloor$ players and scale down the costs of the edges used by the Nash equilibrium by a $f(1, W) + \epsilon$ factor. This yields an (unweighted) game in which the only $f(1, W)$-approximate Nash equilibrium has cost $\Omega(\log W/f(1, W))$ (and the minimum-cost solution still has value 1). Thus $f(1, W) \cdot g(1, W) = \Omega(\log W)$ for all $W \geq 1$. ∎

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 [21]). 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 [6], Rosenthal [19], and Monderer and Shapley [15], 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. [1] 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 [1]). We nonetheless show that ideas from potential functions can be used to derive an essentially 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. On the surface, this function therefore seems incapable of proving an $o(\log W)$ bound on 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. Specifically, 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. Roughly, this allows us to prove that every sequence of such deviations ends sufficiently quickly, without accruing much additional cost. Precisely, we use these ideas to prove the following result (cf., Theorem 4.1). Details are in the Appendix.

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

(a)

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

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.

Remark 5.2 Our proof of Theorem 5.1 is quite flexible and carries over to many of the extensions known for the unweighted case [1]. For example, Theorem 5.1 continues to 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. We defer further details to the full version.

References

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

[2] 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.

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

[4] 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.

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

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

[7] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. Manuscript, 2005.

[8] 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.

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

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

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

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

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

[18] 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.

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

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

[21] T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006. To appear.

[22] 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.

A Missing Proofs

A.1 Missing Proofs from Section 4

A.1.1 Proof of Theorem 4.2

Proof of Theorem 4.2: We prove assertions (1)–(5) in the proof sketch in turn. Before beginning, note that the only (simple) path available to each of A^* and $\bar{A}^*$ is its short path; hence these two players use the edges e_i and $\bar{e}_i$ but no other edges on the primary paths.

For (1), first note that if a small player uses an edge e_j or $\bar{e}_j$ for some $j \in \{1, 2, \ldots, i-1\}$, then it also uses either e_1 or $\bar{e}_1$. We therefore need only prove (1) for $j = 1$. Fix a small player $A_{0,h}$. If this player uses a short path (a path that uses only one edge from each primary path), it incurs a cost of at most

$$1 + \frac{1}{1+p^i}\frac{p^i}{2} \leq \frac{3}{2},$$

since while it might pay for the entire (unit) cost of the edge $e_{0,h}$ or $\bar{e}_{0,h}$, it shares the edge e_i or $\bar{e}_i$ with a player of weight p^i (A^* or $\bar{A}^*$, respectively). Suppose that player $A_{0,h}$ instead uses a path that contains e_1 or $\bar{e}_1$. Then the cost incurred by the player is at least the cost of this edge divided by the sum W of all of the player weights (recall that a small player has unit weight). Our parameter choices ensure that this edge's cost is at least $4p^{i+1}$, while $W = p^{i+1} + 3p^i + \sum_{j=1}^{i-1} p^j$. Since

$$4p^{i+1} \geq (4p - 16)\left(p^i + 3p^{i-1} + \sum_{j=0}^{i-2} p^j\right),$$

the cost incurred by the player on this path is at least $4 - \frac{16}{p} = 4 - 2\epsilon_0$. This is strictly greater than $(2 - \epsilon_0) \cdot \frac{3}{2}$, which establishes (1).

For (2), suppose first that player A_i uses a long path. Such a path must include edge e_1 or $\bar{e}_1$. By part (1), the total weight of the players using this edge is at most $p^i + p^{i-1} + \cdots + p$. Hence the cost incurred by player A_i on such a path is at least

$$4p^{i+1}\frac{p^i}{p^i + p^{i-1} + \cdots + p},$$

which is at least $2p^{i+1}$ since we have chosen p sufficiently large. On the other hand, if A_i chooses a short path, its cost is less than $\frac{3}{4}p^i < 2p^{i+1}/(2 - \epsilon_0)$.

Assertion (3) requires the most involved argument. Suppose A_i uses its lower short path (the argument for the other case is symmetric). First consider player A_{i-1}. If it uses its lower short path, then it shares the first edge (e_{i-1}) with player A_i and hence incurs cost at most

$$\frac{p^{i-1}}{p^{i-1}+p^i} \cdot \frac{p^i}{8} + \frac{p^i}{8}(1+\epsilon) \leq \frac{p^i}{8}\left(\frac{1}{p}+1+\epsilon\right) = \frac{p^i}{8}\left(1+\frac{\epsilon_0}{4}\right).$$

If player A_{i-1} uses any other path, it must in particular use the edges $\bar{e}_{i-1}$ and $\bar{e}_{i-2}$. Moreover, parts (1) and (2) imply that the only other players that could be using these edges are $A_{i-2}, \ldots, A_1$. The cost incurred by A_{i-1} on such a path is therefore at least

$$\frac{p^{i-1}}{p^{i-1}+\cdots+p}\left(\frac{p^i}{8}+\frac{p^i}{8}(1+\epsilon)\right) \geq \frac{p-1}{p} \cdot \frac{p^i}{8} \cdot (2+\epsilon) = \frac{p^i}{8} \cdot \left(1-\frac{\epsilon_0}{8}\right)\left(2+\frac{\epsilon_0}{8}\right).$$

Since

$$\left(1-\frac{\epsilon_0}{8}\right)\left(2+\frac{\epsilon_0}{8}\right) > (2-\epsilon_0)\left(1+\frac{\epsilon_0}{4}\right),$$

player A_{i-1} will choose its lower short path in every $(2-\epsilon_0)$-approximate Nash equilibrium in which player A_i chooses its lower short path.

The above argument then applies inductively to players $A_{i-2}, \ldots, A_2$. For a generic player j, its cost on its lower short path (given that A_{j+1} uses its lower short path) is at most

$$\frac{p^j}{p^j+p^{j+1}} \cdot \frac{p^i}{8}(1+\epsilon)^{i-j-1} + \frac{p^i}{8}(1+\epsilon)^{i-j} < \frac{p^i}{8} \cdot (1+\epsilon)^{i-j-1} \cdot \left(1+\frac{\epsilon_0}{4}\right),$$

while its cost on every other path (given that players $A_{i-1}, \ldots, A_{j+1}$ use their lower short paths) is at least

$$\frac{p^j}{p^j+\cdots+p}\left(\frac{p^i}{8}(1+\epsilon)^{i-j-1}+\frac{p^i}{8}(1+\epsilon)^{i-j}\right) \geq \frac{p^i}{8} \cdot (1+\epsilon)^{i-j-1} \cdot \left(1-\frac{\epsilon_0}{8}\right)\left(2+\frac{\epsilon_0}{8}\right).$$

As above, this implies that A_j will use its lower short path.

Finally, for player A_1, the cost of its lower short path (given that A_2 uses its lower short path) is at most

$$\frac{p}{p+p^2} \cdot \frac{p^i}{8}(1+\epsilon)^{i-2} + n \leq p^{i+1} + 4p^i(1+\epsilon).$$

Every other s_1-t_1 path contains the edge $\bar{e}_1$ which, as a consequence of the previous steps, is otherwise unoccupied. Hence, every other s_1-t_1 path has cost at least $4p^{i+1}$. Since we have chosen p sufficiently large, this is strictly greater than $(2-\epsilon_0)(p^{i+1}+4p^i(1+\epsilon))$, and hence A_1 will use its lower short path, completing the proof of (3).

For (4), first consider the small player $A_{0,1}$. Since parts (1)–(3) imply that edge $\bar{e}_{0,1}$ is unoccupied except possibly for $A_{0,1}$, the player will incur a cost of 1 for using a path that includes this edge. On the other hand, if the player uses its lower short path (the path containing only $e_{0,1}$, $\bar{e}_i$, and zero-cost edges), it shares edge $e_{0,1}$ with player A_1 and therefore incurs cost at most

$$\left(\frac{1}{1+p}\right) \cdot 1 + \left(\frac{1}{1+p^i}\right) \cdot \frac{p^i}{2} \leq \frac{1}{1+p} + \frac{1}{2}.$$

Since $(\frac{1}{2}+\frac{1}{1+(8/\epsilon_0)})(2-\epsilon_0) < 1$, player $A_{0,1}$ will not take a path that includes the edge $\bar{e}_{0,1}$, and must therefore take a path that includes $e_{0,1}$ and wraps around to include the edge $\bar{e}_i$. Applying this same argument inductively to the players $A_{0,2}, \ldots, A_{0,n}$ then proves (4).

Finally, to show (5), consider a $(2-\epsilon_0)$-approximate Nash equilibrium in which all of the small players choose paths that include the edge $\bar{e}_i$ (as usual, the other case is symmetric). If player A_i is using its lower short path, then parts (1)–(4) imply that it shares edge e_i only with player A^*, and therefore its cost share for this edge is $p^i/4$. On the other hand, if A_i uses its upper short path, it shares edge $\bar{e}_i$ with all p^{i+1} of the small players, and therefore its cost on this path is at most

$$\left(\frac{p^i}{p^i+p^{i+1}}\right)\cdot\frac{p^i}{2}+\frac{p^i}{8}\le\frac{p^i}{8}\left(1+\frac{4}{p}\right).$$

Since $\epsilon_0 = 8/p$,

$$\frac{p^i}{4}>\frac{p^i}{8}\left(1+\frac{4}{p}\right)\cdot(2-\epsilon_0),$$

and hence A_i must use its upper short path. This completes the proof of (5) and of the theorem. ■

A.1.2 Proof of Theorem 4.3

In this subsection, we will present an example without $(\frac{c\log W_{max}}{\log\log W_{max}})$-approximate Nash equilibrium for some constant c.

Similar to the construction for Theorem 4.2, this network will consist of $i-1$ stages connected serially. All of the stages except the first one and the last one have the structure shown in Figure 3(a). The first stage has the structure shown in Figure 3(b) and the last stage has the structure shown in Figure 3(c). The cost of the edges are defined as the follows:

- $c(E_{2i}) = c(\bar{E}_{2i}) = p^{2i}$,
- $c(E_{2i-2}) = c(\bar{E}_{2i-2}) = p^{2i}\frac{1}{H(\sqrt{p})}$,
- $c(E_{2i-1}) = c(\bar{E}_{2i-1}) = 3p^{2i}(H(\sqrt{P}))^3$,
- $c(E_{2k,s}) = c(\bar{E}_{2k,s}) = 2^{i-k-1}p^{2i}\frac{1}{s[H(\sqrt{p})]^2}$ for $k=1,2,\ldots,i-2,\ s=1,2,\ldots,\sqrt{p}$,
- $c(E_{2k-1}) = c(\bar{E}_{2k-1}) = 2^{i-k-1}p^{2i+\frac{1}{2}}$ for $k=2,3,\ldots,i-2$,
- $c(E_{0,k}) = c(\bar{E}_{0,k}) = H(\sqrt{p})$, for $k=1,2,\ldots,n$

Here, we set p to be the square of a large integer, $i=\lceil 2.5\log_2 H(\sqrt{p})\rceil+1,\ L=\frac{p^{2i}}{2H(\sqrt{p})}$, $n=2p^{2i}H(\sqrt{p})$.

The players in this network game are the following:

- There are 3 players $P_i,\ P_i^*,\ \bar{P}_i^*$ with weight p^{2i}. They want to Connect from S_i to T_i, from S^* to T^* and from $\bar{S}^*$ to $\bar{T}^*$ respectively.
- There is one player P_4 with weight p^4 who wants to connect from S_4 to T_4, and one player P_2 with weight p^2 who wants to connect from S_2 to T_2.
- There are $\sqrt{p}$ players, $P_{2j,1},\ P_{2j,1},\ \ldots,\ P_{2j,\sqrt{p}}$, with weight p^{2j} for each $j=3,4,\ldots,i-1$. The player $P_{2j,k}$ wants to connect from S_{2j} to $T_{2j,k}$ for every $j,\ k$.
- There are two players P_{2j+1} and $\bar{P}_{2j+1}$ with weight p^{2j+1}, for each $j=1,2,\ldots,i-1$. They wants to connect from S_{2j+1} to T_{2j+1} and from $\bar{S}_{2j+1}$ to $\bar{T}_{2j+1}$, respectively.

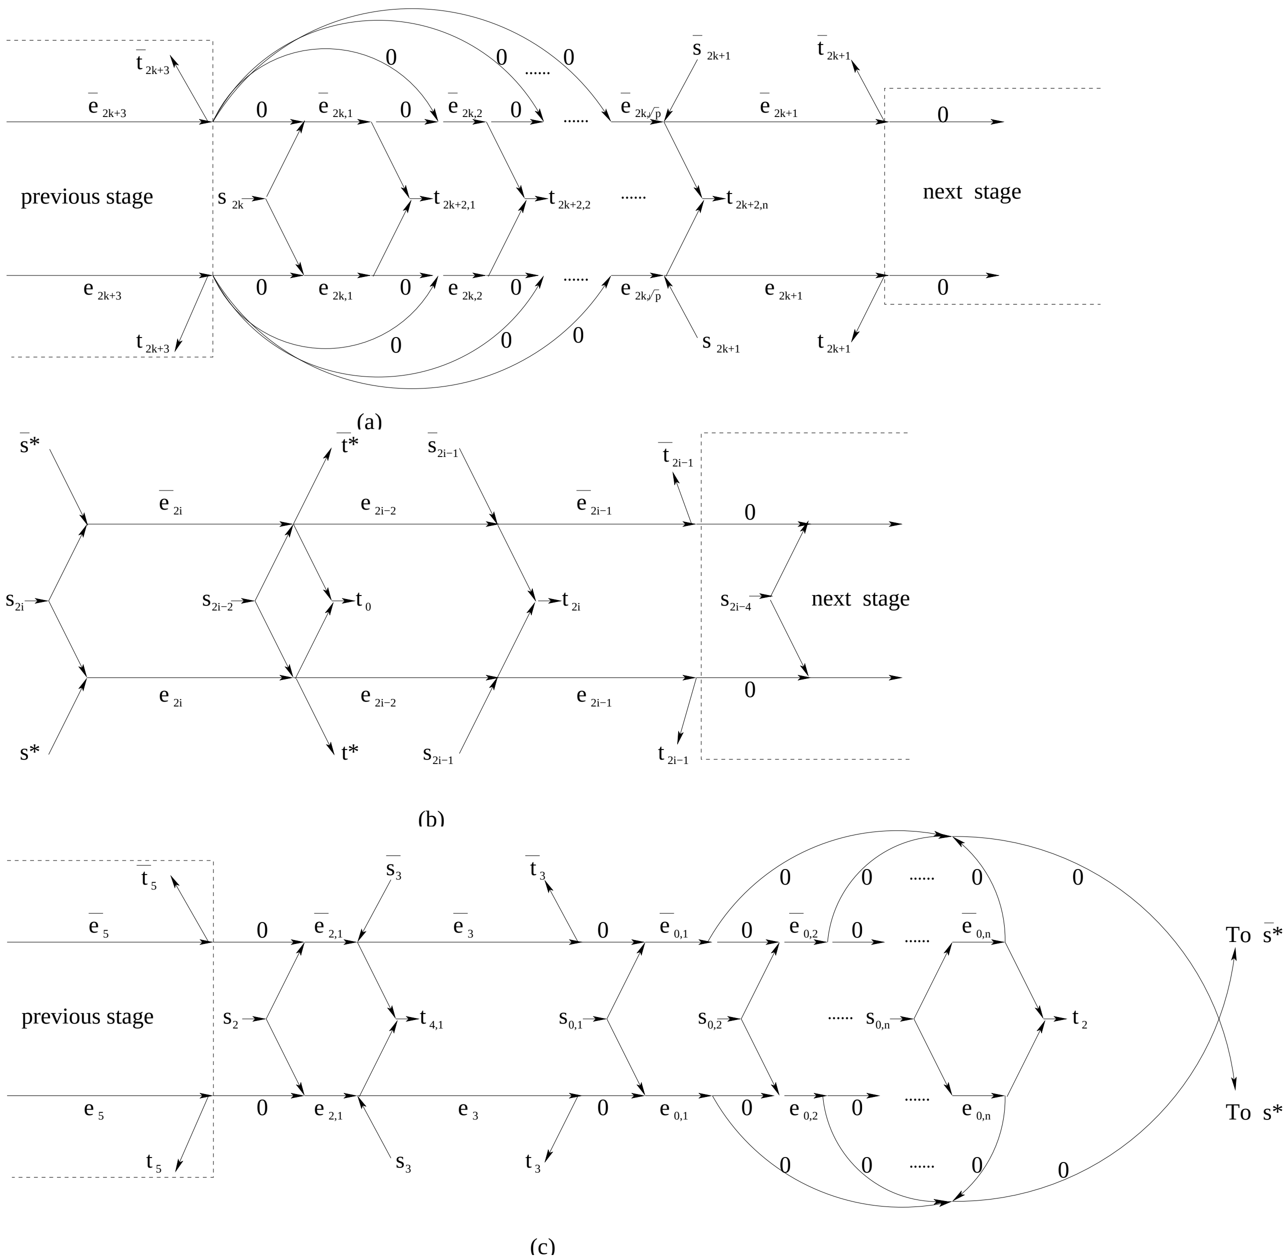

Figure 3: (a) The structure of the $(i-k)$-th stage. (b) The structure of the first stage. (c) The structure of the last stage.

- There are n players $P_{0,1},\ P_{0,2},\ \ldots,\ P_{0,n}$ with weight 1. The k-th player $P_{0,k}$ wants to connect from $S_{0,k}$ to T_0 for every k.

Lemma A.1 *If any player besides P_2 uses some edges in three different stages, then some of the players can deviate from his strategy and reduce the cost by a factor of $H(\sqrt{p})/2$.*

Proof:

- For players P_i^* and $\bar{P}_i^*$, there is only one simple path that connects from S^* to T^* and from $\bar{S}^*$ to $\bar{T}^*$. So, they must use that path and the path only use edges in the first stage. The same proof applies to the players P_{2j+1} and $\bar{P}_{2j+1}$.

- If the player $P_{0,s}$ with weight 1 uses edges in three different stages, then he must use either E_{2i-1} or $\bar{E}_{2i-1}$ to reach the second stage, and his payment is at least $c(E_{2i-1})/(\textit{total weight}) \geq [H(\sqrt{p})]^2$. Hence he can deviate to pay for edges $E_{0,s}$ and $\bar{E}_{2i}$ only, paying at most $H(\sqrt{p})+1$. The cost is reduced by more than a factor of $H(\sqrt{p})/2$ by this switching.

- If the player P_{2i} uses edges in three different stages, then he must use either E_{2i-1} or $\bar{E}_{2i-1}$. His payment for any of these edges is at least $p^{2i}[H(\sqrt{p})]^2$. So, he can deviate to pay for edges E_{2i} and E_{2i-2} only, paying at most $2p^{2i}$. The cost is reduced by more than a factor of $[H(\sqrt{p})]^2/2$.

- Assuming that if any of the players with weight larger than p^{2j} use edges in more than three stages, then some of the players can deviate and reduce the cost by a factor of at least $H(\sqrt{p})/2$. Consider the players $P_{2j,1},\ P_{2j,1},\ \ldots,\ P_{2j,\sqrt{p}}$, if any of them use edges in three different stages, then some of them must use either E_{2j-1} or $\bar{E}_{2j-1}$. If any players with weight larger than p^{2j} uses either E_{2j-1} or $\bar{E}_{2j-1}$, then he needs to use edges in at least three different stages, hence some of the players may deviate from his current strategies and reduce the cost by a factor of at least $H(\sqrt{p})/2$. Otherwise, Any player $P_{2j,s}$ who uses either E_{2j-1} or $\bar{E}_{2j-1}$ must pay at least $2^{i-j-1}p^{2i}$ for that edge, and he may deviate to use edges $E_{2j,1},\ E_{2j,2},\ \ldots,\ E_{2j,\sqrt{p}},\ E_{2j+1}$, and $E_{2j-2,s}$, paying at most $2^{i-j-1}p^{2i}\frac{1}{[H(\sqrt{p})]^2}\sum_{s=1}^{\sqrt{p}}\frac{1}{s}+ 2^{i-j-2}p^{2i+\frac{1}{2}} \times \frac{1}{1+p} + 2^{i-j}p^{2i}\frac{1}{s[H(\sqrt{p})]^2}$, less than $2^{i-j}p^{2i}/H(\sqrt{p})$. The cost is reduced by more than a factor of $H(\sqrt{p})/2$. Inductively, if any of the players $P_{2j,1}, P_{2j,1}, \ldots, P_{2j,\sqrt{p}},\ j = 3, 4, , \ldots, i-1$ uses edges in three different stages, then some of the players can deviate from his current strategy and reduce the cost by at least a factor of $H(\sqrt{p})/2$. The same proof also applies to the player P_4

■

Lemma A.2 *There is no ($H(\sqrt{p})/4$)-approximate Nash equilibrium for the example we constructed.*

Proof: Suppose we are given a set of strategy S which is a ($H(\sqrt{p})/4$)-approximate Nash equilibrium. By lemma A.1, all players except P_2 can only use edges in at most two different stages if they follow the strategy set S. The player P_{2i} has only two choices: using edges $E_{2i},\ E_{2i-2}$ or edges $\bar{E}_{2i},\ \bar{E}_{2i-2}$. Without loss of generality, may assume that he uses $E_{2i},\ E_{2i-2}$.

Now, for the players $P_{2i-2,1},\ P_{2i-1,2},\ \ldots,\ P_{2i-2,\sqrt{p}}$, if some of them uses the edge $\bar{E}_{2i-2}$, let the $P_{2i-2,s}$ be the one with the highest second index among them. Since no other players will use edge $\bar{E}_{2i-2}$ (otherwise they must use edge in at least three stages), the total weight of the players using that edge $\bar{E}_{2i-2}$ is at most sp^{2i-2}. So, the player $P_{2i-2,s}$ must pay at least $\frac{p^{2i}}{sH(\sqrt{p})}$, hence he can deviate and only pay for edges $E_{2i-2},\ E_{2i-1}$ and $E_{2i-4,s}$. His new payment is at most

$p^{2i}/(1+p^2) + 3p^{2i}[H(\sqrt{p})]^3/(1+p) + \frac{2p^{2i}}{s[H(\sqrt{p})]^2}$, which is less than $\frac{4p^{2i}}{s[H(\sqrt{p})]^2}$. So, he can reduce his payment by at least a factor of $H(\sqrt{p})/4$ by this deviation, which is impossible since S is a $H(\sqrt{p})/4$-approximate Nash equilibrium. Therefore, all the players $P_{2i-2,s}$ must use the edge E_{2i-2} and thus must use the edge $E_{2i-4,s}$. So, the edge $E_{2i-4,s}$ is used by the player $P_{2i-2,s}$, for $s = 1, 2, \ldots, \sqrt{p}$ and no players with weight p^{2i-2} use the edges $\bar{E}_{2i-4,s}$, $s = 1, 2, \ldots, \sqrt{p}$. Given this, we can apply the same proof to the players with weight p^{2i-4} and show that they must use all the edges $E_{2i-4,s}$ and not use edges $\bar{E}_{2i-4,s}$. Inductively, we can prove the player $P_{2j,s}$ must use the edge $E_{2j-2,s}$ for all $j = 3, 4, \ldots, i-1$ and $s = 1, 2, \ldots, \sqrt{p}$. Also, player P_4 must use the edge E_2 by the same proof.

From the discussion above and lemma A.1, assuming the player P_{2i} uses E_{2i}, E_{2i-2}, then no players other than P_2 will use the edge $\bar{E}_{2,1}$ according to the set of strategies S. If the player P_2 uses $\bar{E}_{2,1}$, then he must pay at least $c(\bar{E}_{2,1}) = \frac{2^{i-2}p^{2i}}{[H(\sqrt{p})]^2}$, which is more than $p^{2i}[H(\sqrt{p})]^3$. He can deviate to use the edge $E_{2,1}$, E_3 and edges $E_{0,1}, E_{0,2}, \ldots, E_{0,n}$, paying at most $c(E_{2,1})/(1+p^2) + c(E_3)/(1+p) + nH(\sqrt{p})$, which is less than $3p^{2i}[H(\sqrt{p})]^2$. His cost is reduced by at least a factor of $H(\sqrt{p})/4$, a contradiction. So, the player P_2 must use edges $E_{2,1}$, E_3 and edges $E_{0,1}$, $E_{0,2}$, $\ldots$, $E_{0,n}$. If the player $P_{0,1}$ uses the edge $\bar{E}_{0,1}$, he must pay for that edge by himself, hence having to pay at least $H(\sqrt{p})$. But this is impossible since he can choose to pay for edges $E_{0,s}$ and $\bar{E}_{2i}$ only, paying at most $H(\sqrt{p})/(1+p^2)+1$. Similarly, no players with weight 1 use the edges $\bar{E}_{0,s}$ and E_{2i} and all of them use the edge $\bar{E}_{2i}$. So, the player P_{2i} need to pay $p^{2i}/2$ for the edge E_{2i}. But he only have to pay at most $\frac{p^{2i}}{2+2H(\sqrt{p})} + \frac{p^{2i}}{H(\sqrt{p})}$, which is less than $\frac{2p^{2i}}{H(\sqrt{p})}$, if he deviate to use edges $\bar{E}_{2i}$, $\bar{E}_{2i-2}$. Hence the player P_{2i} can deviate and reduce the cost by a factor of $H(\sqrt{p})/2$, a contradiction to the original assumption that S is a $(H(\sqrt{p})/4)$-approximate Nash equilibrium. So, this set of strategy S does not exist, as desired. ■

For the example shown in the section, the maximum weight is $W_{max} = p^{2i}$. Since $H(\sqrt{p}) = \frac{\log 2}{10}\frac{\log W_{max}}{\log\log W_{max}}$, by lemma A.2, there is no $\left(\frac{\log 2}{40}\frac{\log W_{max}}{\log\log W_{max}}\right)$-approximate Nash equilibrium.

A.2 Missing Proofs from Section 5

We now prove Theorem 5.1. We first establish some preliminary results.

Fact A.3 *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 of Fact A.3: 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 the logarithm (base 2) of both sides 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 for (a), raising both sides of this inequality to the $y/(x+y)$ power and then taking the logarithm (base 2) of both sides verifies the claimed inequality. ■

We next consider the existence of approximate Nash equilibria without worrying about their cost. Recall that w_{max} and W denote the maximum player weight and the sum of the player weights of a weighted Shapley network design game, respectively, after weights have been scaled so that the minimum player weight is 1.

Lemma A.4 *For every function $f(w_{max}, W)$ satisfying $f(w_{max}, W) \geq \log_2[e(1 + w_{max})]$ 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. To prove the lemma, it suffices to show 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 $\alpha = f(w_{max}, W)$. We will 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

$$\sum_{e \in Q_i} c_e \cdot \frac{w_i}{W_e + w_i} \leq \frac{1}{\alpha} \sum_{e \in P_i} c_e \cdot \frac{w_i}{W_e}, \tag{3}$$

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:

$$\begin{aligned}
\Delta\Phi &= \sum_{e \in Q_i} c_e[\log_2(1 + W_e + w_i) - \log_2(1 + W_e)] - \\
&\qquad \sum_{e \in P_i} c_e[\log_2(1 + W_e) - \log_2(1 + W_e - w_i)] && (4) \\
&< \sum_{e \in Q_i} c_e \left[\log_2[e(1 + w_i)] \cdot \frac{w_i}{W_e + w_i}\right] - \sum_{e \in P_i} c_e \frac{w_i}{W_e} && (5) \\
&\le \log_2[e(1 + w_{max})] \sum_{e \in Q_i} c_e \frac{w_i}{W_e + w_i} - \sum_{e \in P_i} c_e \frac{w_i}{W_e} \\
&\le -\sum_{e \in P_i} c_e \frac{w_i}{W_e} \times \frac{f(w_{max}, W) - \log_2[e(1 + w_{max})]}{f(w_{max}, W)} && (6) \\
&\le 0.
\end{aligned}$$

In this derivation, the equality (4) follows from the definition of Φ; the inequality (5) follows from Fact A.3, with Fact A.3(b) applied to each term in the first sum with $x = W_e$ and $y = w_i$, and Fact A.3(a) applied to each term in the second sum with $x = W_e - w_i$ and $y = w_i$; and the final inequality (6) follows from (3) and our choice of α. ■

We now extend the argument in the proof of Lemma A.4 to account for the cost of approximate equilibria, which proves Theorem 5.1.

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 A.4, 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

$$\begin{aligned}
\Delta\Phi &\le -\sum_{e \in P_i} c_e \frac{w_i}{W_e} \times \frac{f(w_{max}, W) - \log_2[e(1 + w_{max})]}{f(w_{max}, W)} && (7) \\
&\le -\frac{1}{2} \sum_{e \in P_i} c_e \frac{w_i}{W_e} && (8) \\
&\le -\frac{1}{2} A \times f(w_{max}, W), && (9)
\end{aligned}$$

where inequality (7) is the same as inequality (6) in the proof of Lemma A.4; inequality (8) follows from the choice of the function f; and inequality (9) 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 social cost increases by an additive factor of A, the potential function Φ decreases by at least $\frac{1}{2} f(w_{max}, W) \cdot A$. The potential function value of the social optimum is at most a $\log_2(1 + W)$ multiplicative factor larger than its cost C^*, and the potential function only decreases throughout the sequence of deviations. The social cost can therefore only increase by a $2C^* \log_2(1 + W)/f(w_{max}, W)$ additive factor throughout the entire sequence of deviations. The sequence must therefore terminate in a $\left(f(w_{max}, W), 1 + \frac{2\log_2(1+W)}{f(w_{max},W)}\right)$-approximate Nash equilibrium. ■