

Improving the H_k -Bound on the Price of Stability in Undirected Shapley Network Design Games

Yann Disser¹, Andreas Emil Feldmann², Max Klimm¹, and Matúš Mihalák³

¹ Institut für Mathematik, Technische Universität Berlin, Germany
{disser,klimm}@math.tu-berlin.de

² Combinatorics and Optimization Department, University of Waterloo, Canada
andreas.feldmann@uwaterloo.ca

³ Institute of Theoretical Computer Science, ETH Zurich, Switzerland
matus.mihalak@inf.ethz.ch

Abstract. In this paper we show that the price of stability of Shapley network design games on undirected graphs with k players is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2}H_k = (1 - \Theta(1/k^4))H_k$, where H_k denotes the k -th harmonic number. This improves on the known upper bound of H_k , which is also valid for directed graphs but for these, in contrast, is tight. Hence, we give the first non-trivial upper bound on the price of stability for undirected Shapley network design games that is valid for an arbitrary number of players. Our bound is proved by analyzing the price of stability restricted to Nash equilibria that minimize the potential function of the game. We also present a game with $k = 3$ players in which such a restricted price of stability is 1.634. This shows that the analysis of Bilò and Bove (Journal of Interconnection Networks, Volume 12, 2011) is tight. In addition, we give an example for three players that improves the lower bound on the (unrestricted) price of stability to 1.571.

Keywords: undirected Shapley network design game, price of stability, potential-optimal price of stability, potential-optimal price of anarchy.

1 Introduction

Infrastructure networks are the lifelines of our civilization. Through generations a tremendous effort has been undertaken to cover the earth's surface with irrigation canal systems, sewage lines, road networks, railways, and – more recently – data networks. Some of these infrastructures are initiated and planned by a central authority that designs the network and decides on its topology and dimension. Many networks, however, arise as an outcome of actions of selfish individuals who are motivated by their own connectivity requirements rather than by optimizing the overall network design. A prominent example of the latter phenomenon is the rise of the Internet. In order to quantify the efficiency of networks, it is crucial to understand the processes that govern their formation. Anshelevich et al. [1] proposed a particularly elegant model for such processes, which is now known as the Shapley network design game or the network design game with fair cost allocation (for an overview of other models for network formation, see [2]).

The Shapley network design game is played by k players $1, 2, \dots, k$ on a graph $G = (V, E)$ with positive edge-costs $c_e \in \mathbb{N}$. Each player i has associated with it a source-target pair $s_i, t_i \in V$ of vertices that she needs to connect with a simple path in G . The choice of such a path is called a *strategy* of the player, and a collection consisting of one strategy for each player is called a *strategy profile*. The cost c_e of every edge e is shared equally among the players using it. Each player i aims at choosing a path of smallest possible (individual) cost to herself. This cost is defined as the sum of the cost shares for player i along the path. Players are selfish in that they only care about their own costs. In particular, they do not care about the social cost, defined as the sum of all players' individual costs and denoted by $\text{cost}(P)$ for a strategy profile P .

A *Nash equilibrium* of a Shapley network design game is a strategy profile in which no player i can switch to an s_i - t_i path that yields her a smaller individual cost. To quantify the effect of the selfish behavior, it is natural to compare the social cost of a Nash equilibrium of the game with the smallest social cost among all possible strategy profiles [1,2]. Several quantifications of selfish behavior have been studied, based on whether we restrict ourselves to a specific set of Nash equilibria, and whether we compare the worst or best such equilibrium in terms of social cost. In this paper, we adopt the notion of the *price of stability*, introduced by Anshelevich et al. [1]. Denoting by \mathcal{N} the set of all Nash equilibria and by O a strategy profile that minimizes the social cost of a game, the price of stability of the game is defined as the ratio $\min_{N \in \mathcal{N}} \text{cost}(N) / \text{cost}(O)$.

Anshelevich et al. [1] observed that Shapley network design games always have a Nash equilibrium by showing that they belong to the class of congestion games. For these games, the existence of a Nash equilibrium is always guaranteed, as shown by Rosenthal [3]. Rosenthal's existence proof relies on a potential function argument. That is, he showed that there exists a function Φ that maps strategy profiles to real numbers and has the property that if any one player changes her strategy unilaterally, then the value of Φ changes by the exact same value as the cost of the player. This observation, together with the finiteness of the space of all strategy profiles, implies the existence of a Nash equilibrium. In particular, any *potential minimum*, i.e., a strategy profile that globally minimizes the potential function, is a Nash equilibrium. The potential function of a game is unique up to an additive constant (see Monderer and Shapley [4]). Using the special form of the potential function for Shapley network design games, Anshelevich et al. [1] showed that the price of stability of any game is at most $H_k = \sum_{i=1}^k \frac{1}{i}$, the k -th *harmonic number* (which is of order $\log k$). This upper bound is tight for games played on directed graphs. That is, there are Shapley network design games on directed graphs [1] for which the price of stability is arbitrarily close to H_k .

The situation is different for *undirected* Shapley network design games, i.e., games played on undirected graphs. As the same potential arguments remain valid, the price of stability of any game is still at most H_k . Yet, the largest known price of stability (asymptotically) is a constant, more precisely $348/155 \approx 2.245$ (see Bilò et al. [5]). This leaves the question of the worst-possible price of stability in undirected Shapley network design games with k players wide

open. Remarkably, the largest known price of stability, as provided by Bilò et al., does not come from a simple example, but from a complicated construction. Previously known worst-case games had a price of stability of $4/3 \approx 1.333$ [1], $12/7 \approx 1.714$ [6], and $42/23 \approx 1.8261$ [7]. Despite numerous attempts [8,1,5,6,7,9] to narrow the gap of the bounds on the price of stability, there has been little progress in terms of numerical results. It is generally believed that the price of stability is smaller than H_k , and we confirm this belief in this paper. For small values of k some smaller upper bounds are known. For $k = 2$ players, the price of stability is at most $4/3 < H_2 = 3/2$ and this is tight [1,7]. Bilò and Bove [8] analyzed the case of $k = 3$ players and showed that the price of stability of any such game is at most $1.634 < H_3 = 1.83\bar{3}$. For this case, however, a considerable gap remains, as the worst example known has a price of stability of $74/48 \approx 1.542$ [7]. Thus, already for $k = 3$ players, the exact worst-case price of stability is unknown.

For several special cases, one can derive better upper bounds on the price of stability. If all players share the same terminal then the price of stability is at most $O(\log k / \log \log k)$ [9]. If in addition every vertex is the source of at least one player, then the price of stability further degrades to $O(\log \log k)$ [6].

Many of the mentioned upper bounds are not only valid for the best Nash equilibrium of a game, but also for a very specific one – the potential minimum. Potential minima have desirable stability properties. For example, they are reached by certain learning dynamics for players that do not always play rationally (see Blume [10]). This motivates to explicitly study the ratio between the cost of a potential minimum and that of a profile minimizing the social cost – a *social optimum*. To stress the described stability properties of potential minima, Asadpour and Saberi [11] called this ratio the *inefficiency ratio of stable equilibria*. Kawase and Makino [12] called the very same ratio the *potential-optimal price of anarchy*. They also define the *potential-optimal price of stability* of a game in the obvious way as the ratio between the cost of a best potential minimum and that of a social optimum. They prove that the potential-optimal price of anarchy of undirected Shapley network design games is at most $O(\sqrt{\log k})$ for the special case where all players share the same terminal node, and where every vertex is the source of at least one player. They give a construction of a game with potential-optimal price of anarchy $\Omega(\sqrt{\log \log k})$.

Our Contribution

Our main result shows that the price of stability in undirected Shapley network design games is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k = (1 - \Theta(1/k^4))H_k$. Thus, we provide the first general upper bound that shows that the price of stability for k players is strictly smaller than H_k . To prove this upper bound, we generalize the techniques of Christodoulou et al. [7] to any number of players. In short, similar to Christodoulou et al., we obtain a set of inequalities relating the cost of any Nash equilibrium to the cost of a social optimum. We then combine these in a non-trivial way to obtain the claimed upper bound, additionally assuming that the

Nash equilibrium has a smaller potential than the social optimum. Interestingly, the resulting upper bound is tight for the case of $k = 2$ players.

As an additional contribution, we provide an example of a game with $k = 3$ players in which the potential-optimal price of stability is 1.634. Thus, we show that the upper bound on the potential-optimal price of anarchy given by Bilò and Bove [8] is tight. This result implies that for three players the upper bound on the price of stability cannot be further improved via potential-minimizers. This is in contrast to the directed case, for which a simple inequality relates the cost of the potential minimum to that of a social optimum, giving a tight bound on the price of stability. We believe that this observation provides an insight as to why the undirected case is much harder to tackle than the directed one.

We note that in our tight example for the potential-optimal price of stability/anarchy, the social optimum is also a Nash equilibrium, and thus the example provides no new lower bounds on the price of stability. Our third contribution however is a new lower bound on the price of stability. We provide an example of a game with three players and price of stability 1.571, which improves on the previous best lower bound of $74/48 \approx 1.542$ [7].

2 Problem Definition and Preliminaries

Let $G = (V, E)$ be an undirected graph with a positive *cost* $c_e > 0$ for every edge $e \in E$. The *Shapley network design game* is a strategic game of k players. Every player $i \in \{1, \dots, k\}$ has a dedicated pair of vertices $s_i, t_i \in V$, that we call her *source* and *target*, respectively. The *strategy space* of player i is the collection \mathcal{P}_i of all paths $P_i \subseteq E$ between s_i and t_i . Every such path P_i is called a *strategy*. A *strategy profile* P is a tuple (P_1, \dots, P_k) of k strategies, $P_i \in \mathcal{P}_i$. Given a strategy profile P , we say that player i *plays* strategy P_i in P . The *cost to player i* in a strategy profile $P = (P_1, \dots, P_k)$ is $\text{cost}_i(P) = \sum_{e \in P_i} \frac{c_e}{k_e}$, where k_e denotes the number of paths P_i in P such that $e \in P_i$. That is, k_e is the number of players that use edge e . The goal of every player is to minimize her cost. A *Nash equilibrium* is a strategy profile $N = (N_1, \dots, N_k)$, $N_i \in \mathcal{P}_i$, such that no player i can improve her cost by playing a different strategy. That is, for every i and every $N'_i \in \mathcal{P}_i$, it holds that $\text{cost}_i(N) \leq \text{cost}_i(N'_i, N_{-i})$, where (N'_i, N_{-i}) is a shorthand for $(N_1, \dots, N_{i-1}, N'_i, N_{i+1}, \dots, N_k)$. With a slight abuse of terminology we will identify the game played on the graph G with the graph itself and we write $\mathcal{N}(G)$ to denote the set of Nash equilibria of G .

Observe that the edges of any strategy profile P induce a graph $(V, \cup_i P_i)$, which we call the *underlying network*. We denote the edge set of this graph by $E(P)$. The *social cost*, or simply the *cost* of a strategy profile P , denoted as $\text{cost}(P)$, is the sum of the players' individual costs. Observe that the social cost is equal to the total cost of the edges in the played strategies, i.e., $\text{cost}(P) = \sum_{i=1}^k \text{cost}_i(P) = \sum_{e \in E(P)} c_e$.

A strategy profile that minimizes the social cost is called the *social optimum*. The *price of stability* of a game G , denoted by $\text{PoS}(G)$, is defined as the cost of the best Nash equilibrium of G divided by the cost of a social optimum $O(G)$

of G . That is, $\text{PoS}(G) = \min_{N \in \mathcal{N}(G)} \text{cost}(N)/\text{cost}(\text{O}(G))$. The *price of anarchy* of G , $\text{PoA}(G)$ for short, is obtained by replacing \min by \max in this definition.

A Shapley network design game is a potential game [1,3]. That is, there is a function $\Phi : \mathcal{P}_1 \times \dots \times \mathcal{P}_k \rightarrow \mathbb{R}$ such that, for every strategy profile P , whenever any player i changes her strategy from P_i to P'_i , then $\text{cost}_i(P) - \text{cost}_i(P'_i, P_{-i}) = \Phi(P) - \Phi(P'_i, P_{-i})$. Rosenthal's (exact) potential function has the form $\Phi(P) = \sum_{e \in E(P)} \sum_{i=1}^{k_e} \frac{c_e}{i} = \sum_{e \in E(P)} H_{k_e} \cdot c_e$, where H_j denotes the j -th Harmonic number $\sum_{i=1}^j \frac{1}{i}$. The potential function is unique up to an additive constant [4].

Motivated by the particular stability properties of potential minima, Kawase and Makino [12] introduced two notions to quantify the inefficiency of potential minimizers. For a game G let $\mathcal{F}(G)$ denote the set of potential minimizers of G , i.e., strategy profiles of G that minimize the potential function of G . The *potential-optimal price of stability* of G is then defined as $\text{POPoS}(G) = \min_{N \in \mathcal{F}(G)} \text{cost}(N)/\text{cost}(\text{O}(G))$ and the *potential-optimal price of anarchy* is defined as $\text{POPoA}(G) = \max_{N \in \mathcal{F}(G)} \text{cost}(N)/\text{cost}(\text{O}(G))$. Since $\mathcal{F}(G) \subseteq \mathcal{N}(G)$, clearly, for any game G , $\text{PoS}(G) \leq \text{POPoS}(G) \leq \text{POPoA}(G) \leq \text{PoA}(G)$.

For a fixed number of players $k \geq 2$, we are interested to bound the *worst-case* price of stability of games with k players. For a formal definition, let $\mathcal{G}(k)$ denote the set of all games with k players. The price of stability of undirected Shapley network design games with k players is defined as $\text{PoS}(k) = \sup_{G \in \mathcal{G}(k)} \text{PoS}(G)$. $\text{POPoS}(k)$ and $\text{POPoA}(k)$ are defined analogously.

Using Rosenthal's potential function Φ , we can bound the potential-optimal price of anarchy (and, thus, the price of stability) from above by H_k as follows (cf. [1]). Let O be a social optimum. For a potential minimum N , we have $\Phi(N) \leq \Phi(\text{O})$. Using this together with $\text{cost}(N) \leq \Phi(N)$ and $\Phi(\text{O}) \leq H_k \cdot \text{cost}(\text{O})$ we obtain $\text{cost}(N) \leq H_k \cdot \text{cost}(\text{O})$, as claimed.

We define N^i and O^i , for $i \in \{1, \dots, k\}$, to be the sets of edges of N and O that are used by exactly i players, respectively. Thus, $E(N) = \bigcup_{i=1}^k N^i$ and $E(O) = \bigcup_{i=1}^k O^i$. For a set of edges $M \subseteq E$, we will denote by $|M|$ the total cost of the edges in M , i.e., $|M| = \sum_{e \in M} c_e$. This allows us to express the value of the potential function for N and O by $\Phi(N) = \sum_{j=1}^k H_j |N^j|$ and $\Phi(O) = \sum_{j=1}^k H_j |O^j|$, respectively.

3 A General Upper Bound

In this section we derive an upper bound on the price of stability of *any* undirected Shapley network design game. The main idea to show our upper bound follows that of Christodoulou et al. [7], which they used for deriving an upper bound of $33/20 = 1.65$ for the case of $k = 3$ players. By the definition of Nash equilibria, no player can change her chosen s_i - t_i path in such a profile and thereby improve her cost. We will consider a specific change of strategy for each player, which gives us an inequality that relates the costs of the edges used in the Nash equilibrium to those used in the social optimum. We then combine this inequality in a non-trivial way with another that is gained from the fact that

we consider a potential minimum. This allows us to obtain an upper bound of $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2}H_k = (1 - \Theta(1/k^4))H_k$ on the price of stability.

Consider a Nash equilibrium N and a social optimum O of a given Shapley network design game with k players. In undirected graphs, if O^k is non-empty any player i can use paths in the underlying network of the optimum O to connect its terminals to the source and the target of another player j . By additionally using the path of player j in the Nash equilibrium N , we obtain a valid strategy for player i . This is the specific alternative strategy of player i which we use to relate the cost of the Nash equilibrium to the social optimum.

We will additionally use the following observation about the structure of social optima. Observe that any social optimum O forms a forest (because from any cycle we could remove an edge and thereby decrease the cost). If O^k is non-empty, i.e., some edges are shared among all players in the optimum, the edges of $E(O) \setminus O^k$ form two trees such that every player has one terminal in each of the trees. Let O^+ denote the edge set of the larger tree, and O^- denote the edge set of the smaller one (measured in terms of the social cost). We have $O^+ \cup O^- = E(O) \setminus O^k$ and, by the definition, $|O^+| \geq |O^-|$. Every tree has a closed walk that visits every vertex at least once and every edge exactly twice (for example, a depth-first traversal). We consider such a walk in the tree given by the edges in O^+ , and use it to order the players. Without loss of generality, if $O^k \neq \emptyset$, the players are numbered such that there is a closed walk in O^+ that visits the terminals in the order given by the players' numbers, while using each edge exactly twice. We say that the players are in *major-tree order*.

Consider the edges of a Nash equilibrium which are not used by all k players. The following lemma bounds the cost of these edges with respect to the cost of a social optimum.

Lemma 1. *Given a game G , let $N = (N_1, \dots, N_k)$ be a Nash equilibrium and $O = (O_1, \dots, O_k)$ a social optimum with $O^k \neq \emptyset$. Then,*

$$\sum_{j=1}^{k-1} |N^j| \leq (k^2(k+1)/2 - k) \sum_{j=1}^{k-1} |O^j|.$$

Proof. For every player i , we construct an s_i - t_i path P_i (cf. Figure 1) with the property that every edge on P_i is either in N_{i+1} or not in O^k . In the following we understand indices modulo k , i.e. $k+1 \equiv 1$ and $0 \equiv k$. Let u, v be the first and last vertex on O_i that are also on O_{i+1} . These vertices are well defined as O^k is non-empty, and thus $O_i \cap O_{i+1}$ is non-empty. Assume that u lies before v on O_{i+1} (otherwise, exchange s_{i+1} and t_{i+1} in the following). Let P_i be the path from s_i to t_i that first follows O_i until u , then O_{i+1} (backwards) from u to s_{i+1} , then N_{i+1} from s_{i+1} to t_{i+1} , then O_{i+1} (backwards) from t_{i+1} to v , and finally O_i from v to t_i . In case P_i contains cycles, we skip them to obtain a simple path. It is easy to verify that every edge on P_i either lies on N_{i+1} or is not in $O_i \cap O_{i+1}$. Thus, P_i has the desired property. In the following, let Q_i denote the set of edges of P_i that are also contained in $E(O)$.




Fig. 1. Constructing the path P_i with O_i, O_{i+1}, N_{i+1} (dashed line), and Q_i (continuous lines). Note that $O_i \cap O_{i+1}$ (dotted line) is not part of P_i .

Since N is a Nash equilibrium, player i cannot improve her cost by choosing path P_i . Therefore, $\text{cost}_i(N) \leq \text{cost}_i(P_i, N_{-i})$. If k_e is the number of players using edge e in N , this inequality amounts to $\sum_{e \in N_i} \frac{c_e}{k_e} \leq \sum_{e \in P_i \cap N_i} \frac{c_e}{k_e} + \sum_{e \in P_i \setminus N_i} \frac{c_e}{k_e + 1}$. By the properties of P_i , the right-hand side of this inequality can be upper bounded by $\sum_{e \in Q_i} c_e + \sum_{e \in N_{i+1} \cap N_i} \frac{c_e}{k_e} + \sum_{e \in N_{i+1} \setminus N_i} \frac{c_e}{k_e + 1}$. By shifting all terms not depending on Q_i to the left-hand side of the resulting inequality, we get

$$\sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e} - \sum_{e \in N_{i+1} \setminus N_i} \frac{c_e}{k_e + 1} \leq \sum_{e \in Q_i} c_e. \tag{1}$$

Similar to the path P_i we can define a path \hat{P}_i with respect to player $i - 1$. That is, \hat{P}_i uses the edges of O_i, O_{i-1} , and N_{i-1} to connect s_i to t_i and does not contain any edges from $O_i \cap O_{i-1}$. Let \hat{Q}_i denote the set of edges of \hat{P}_i also contained in $E(O)$. Using the same arguments as above on \hat{P}_i we can derive an analogous inequality as (1) for edges in N_i, N_{i-1} , and \hat{Q}_i . Adding this inequality to (1) and then summing over all i gives

$$\sum_{i=1}^k \left(\sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e} - \sum_{e \in N_{i+1} \setminus N_i} \frac{c_e}{k_e + 1} + \sum_{e \in N_i \setminus N_{i-1}} \frac{c_e}{k_e} - \sum_{e \in N_{i-1} \setminus N_i} \frac{c_e}{k_e + 1} \right) \leq \sum_{i=1}^k \left(\sum_{e \in Q_i} c_e + \sum_{e \in \hat{Q}_i} c_e \right). \tag{2}$$

We bound the left-hand side and the right-hand side of (2) separately, starting with the left-hand side. Since indices are modulo k , we may shift the index in the second and the fourth sum on the left-hand side. This lets us combine the first and the fourth sum to $\sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e} - \sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e + 1} = \sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e(k_e + 1)}$, and analogously the second and the third sum to $\sum_{e \in N_i \setminus N_{i-1}} \frac{c_e}{k_e(k_e + 1)}$, to obtain

$$\sum_{i=1}^k \sum_{e \in N_i \setminus N_{i+1}} \frac{c_e}{k_e(k_e + 1)} + \sum_{i=1}^k \sum_{e \in N_i \setminus N_{i-1}} \frac{c_e}{k_e(k_e + 1)}.$$

Each of the two resulting sums counts each edge in $E(N) \setminus N^k$ at least once. This is because for any edge e used by at least one player but not all of them, there is a pair of players with consecutive indices (modulo k) such that e is used in N by one of the players but not the other. Thus, we can lower bound the above term by

$$2 \sum_{e \in E(N) \setminus N^k} \frac{c_e}{k_e(k_e + 1)} = \sum_{j=1}^{k-1} \frac{2}{j(j+1)} |N^j| \geq \frac{2}{k(k-1)} \sum_{j=1}^{k-1} |N^j|,$$

which gives a lower-bound on the left-hand side of Inequality (2).

The right-hand side of Inequality (2) can be bounded by exploiting the major-tree order of the players. We first only bound the sum depending on Q_i . We denote the two parts of Q_i that lie in the larger and smaller parts of $E(O) \setminus O^k$ by $Q_i^+ = Q_i \cap O^+$ and $Q_i^- = Q_i \cap O^-$, respectively. Note that, by construction of P_i , there are no edges of O^k in Q_i . Thus, we get

$$\sum_{i=1}^k \sum_{e \in Q_i} c_e = \sum_{i=1}^k \left(\sum_{e \in Q_i^+} c_e + \sum_{e \in Q_i^-} c_e \right).$$

By the defining property of the major-tree order, each edge in O^+ is counted exactly twice in the above sum, while each edge of O^- is counted at most k times. At the same time, the weight of the edges in O^- amounts to at most half the total weight of $E(O) \setminus O^k$. Hence,

$$\begin{aligned} \sum_{i=1}^k \left(\sum_{e \in Q_i^+} c_e + \sum_{e \in Q_i^-} c_e \right) &\leq 2 \sum_{e \in O^+} c_e + k \sum_{e \in O^-} c_e = 2 \sum_{e \in E(O) \setminus O^k} c_e + (k-2) \sum_{e \in O^-} c_e \\ &\leq (k/2 + 1) \sum_{e \in E(O) \setminus O^k} c_e. \end{aligned}$$

Analogously, we can derive a corresponding bound for the sum depending on \hat{Q}_i . Since the sum over all costs of edges in $E(O) \setminus O^k$ is exactly $\sum_{i=1}^{k-1} |O^i|$, we can bound the right-hand side of Inequality (2) by $(k+2) \sum_{i=1}^{k-1} |O^i|$. Together the two derived bounds for the left-hand side and the right-hand side of Inequality (2) give the claimed inequality. \square

The following lemma encapsulates some technical calculations that allows to derive an upper bound on the price of stability using Lemma 1.

Lemma 2. *For a game G with social optimum O , let N be a Nash equilibrium with $\Phi(N) \leq \Phi(O)$ and let $\beta > 0$ be such that $\sum_{j=1}^{k-1} |N^j| \leq \beta \sum_{j=1}^{k-1} |O^j|$. Then,*

$$\sum_{j=1}^k |N^j| \leq \frac{\beta k}{1 + \beta k} H_k \cdot \sum_{j=1}^k |O^j|.$$

Proof. We take $1 < \alpha < H_k$ and compute

$$\begin{aligned} \sum_{j=1}^k |N^j| &\leq \alpha |N_k| + \sum_{j=1}^{k-1} |N^j| = \frac{\alpha}{H_k} \cdot \sum_{j=1}^k H_j |N^j| + \sum_{j=1}^{k-1} \left(1 - \alpha \frac{H_j}{H_k}\right) |N^j| \\ &\leq \frac{\alpha}{H_k} \cdot \sum_{j=1}^k H_j |N^j| + \sum_{j=1}^{k-1} \left(1 - \frac{\alpha}{H_k}\right) |N^j|. \end{aligned}$$

Using that $\Phi(N) \leq \Phi(O)$ and the condition of the lemma, we introduce β to get

$$\begin{aligned} \sum_{j=1}^k |N^j| &\leq \frac{\alpha}{H_k} \cdot \sum_{j=1}^k H_j |O^j| + \left(1 - \frac{\alpha}{H_k}\right) \cdot \beta \sum_{j=1}^{k-1} |O^j| \\ &= \alpha |O^k| + \sum_{j=1}^{k-1} \left[\alpha \frac{H_j}{H_k} + \beta \left(1 - \frac{\alpha}{H_k}\right) \right] |O^j| \\ &\leq \alpha |O^k| + \sum_{j=1}^{k-1} \left[\alpha \frac{H_{k-1}}{H_k} + \beta \left(1 - \frac{\alpha}{H_k}\right) \right] |O^j|. \end{aligned} \tag{3}$$

For $\alpha = \frac{\beta k}{1 + \beta k} H_k$ we have

$$\begin{aligned} \alpha \frac{H_{k-1}}{H_k} + \beta \left(1 - \frac{\alpha}{H_k}\right) &= \frac{\beta k}{1 + \beta k} H_{k-1} + \beta \left(1 - \frac{\beta k}{1 + \beta k}\right) \\ &= \frac{\beta k}{1 + \beta k} \left(H_k - \frac{1}{k}\right) + \frac{\beta}{1 + \beta k} = \alpha. \end{aligned} \tag{4}$$

From (3) and (4) it follows that

$$\sum_{j=1}^k |N^j| \leq \alpha |O^k| + \alpha \sum_{j=1}^{k-1} |O^j| = \frac{\beta k}{1 + \beta k} H_k \sum_{j=1}^k |O^j|,$$

which concludes the proof. \square

The above lemmas can be put together in order to show the following theorem.

Theorem 3. *The potential-optimal price of anarchy POPoA(k) for Shapley network design games with $k \geq 2$ players is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k$.*

Proof. Let N be a potential minimum. If $O^k \neq \emptyset$, we may combine Lemma 1 and Lemma 2 to obtain $\frac{\text{cost}(N)}{\text{cost}(O)} \leq \frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k$. If, on the other hand, $O^k = \emptyset$, then obviously also $|O^k| = 0$. We show that then $\text{cost}(N) \leq H_{k-1} \text{cost}(O)$ (which has been observed before for $k = 3$, e.g., by Christodoulou et al. [7]). We can express the potential functions of N and O using N^j and O^j and use the fact that $\Phi(N) \leq \Phi(O)$ to obtain

$$\sum_{j=1}^k |N^j| \leq \sum_{j=1}^k H_j |N^j| \leq \sum_{j=1}^k H_j |O^j| = \sum_{j=1}^{k-1} H_j |O^j| \leq H_{k-1} \sum_{j=1}^k |O^j|.$$

Hence, in this case we get $\frac{\text{cost}(\text{N})}{\text{cost}(\text{O})} \leq H_{k-1}$, which is lower than the first bound. \square

We obtain the following corollary. Note that the bound is tight for $k = 2$.

Corollary 4. *The price of stability $\text{PoS}(k)$ for Shapley network design game with $k \geq 2$ players is at most $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k$.*

4 Three-Player Games

The upper bound of $\frac{k^3(k+1)/2-k^2}{1+k^3(k+1)/2-k^2} H_k$ on the price of stability presented in the previous section is valid for an arbitrary number of players. For the special case of $k = 3$ players, it evaluates to $165/92 \approx 1.793$. For this case, however, better bounds are known. Bilò and Bove proved that the price of stability does not exceed $286/175 \approx 1.634$. For the proof of their result they combine inequalities that are valid for any potential minimum of the game. Thus, their proof implies that also the potential-optimal price of anarchy (and thus the potential-optimal price of stability) is at most $286/175$. As the main result of this section, we will show that this result is tight. That is, there is a three-player game such that the cost of the best potential minimum is $286/175$ times the cost of the social optimum.

Theorem 5. *For three players, the potential-optimal price of stability and the potential-optimal price of anarchy are $\text{POPoA}(3) = \text{POPoS}(3) = 286/175 \approx 1.634$.*

Proof. The upper bound of $286/175$ on the potential-optimal price of anarchy was proved by Bilò and Bove [8, Theorem 3.1]. They derive the bound for the potential-optimal price of anarchy, but only explicitly state the implied bound for the (regular) price of stability.

Consider the three player game in Figure 2(a), and let $\epsilon > 0$ be sufficiently small. The potential-optimal price of stability of this example approaches $286/175$ when ϵ tends to 0, which establishes tightness of the upper bound. It also shows that the potential-optimal price of stability and the potential-optimal price of anarchy coincide for the class of games with three players.

Obviously, any strategy profile in the example has to use at least three edges to connect all terminal pairs. The three cheapest edges already connect all terminal pairs and thus constitute the social optimum O of cost $700 + 3\epsilon$. It is easy to verify that the underlying networks of Nash equilibria in the example do not contain cycles, since at least one edge of each cycle would be abandoned by all players. Hence, all Nash equilibria in the example use exactly three edges. We show that the *unique* potential minimum N uses the three edges not used in O (note that O itself is a Nash equilibrium). This profile has both a potential-function value and cost equal to $396 + 2 \cdot 374 = 1144$, since every edge is used by one player only. In contrast, the social optimum has a potential-function value of $2H_2 \cdot (209 + \epsilon) + H_3 \cdot (282 + \epsilon) > 1144$. If the edge $\{t_2, s_3\}$ is used in a Nash equilibrium, the other two edges have to be used by at least two players each. For profiles other than O , this gives a potential function value of at least




Fig. 2. (a) A three-player game with POPoS and POPa approaching $286/175 \approx 1.634$ for $\epsilon \rightarrow 0$. (b) A three-player game with PoS = $1769/1126 \approx 1.571$.

$H_2 \cdot (209 + \epsilon + 374) + (282 + \epsilon) > 1156$. If the edges with cost 396 and $(282 + \epsilon)$ are both unused, all three players use an edge of cost 374 and the resulting potential value is either $H_3 \cdot 374 + H_2 \cdot (209 + \epsilon) + (209 + \epsilon) > 1208$ or $H_2 \cdot 374 + 374 + (209 + \epsilon) > 1144$. Equilibria that use one edge each of costs $(209 + \epsilon)$, 374, 396 have a potential function value of $H_2 \cdot 396 + 374 + (209 + \epsilon) > 1177$. And finally, the profile using both cheap edges together with the one of cost 396 has a potential of $H_3 \cdot 396 + 2 \cdot (209 + \epsilon) > 1144$. We conclude that the potential minimum is as claimed. For ϵ tending to 0, the ratio between the cost of N and O approaches $286/175$. \square

Our result in particular implies that it is impossible to push the upper bound on the price of stability for three-player Shapley network design games on undirected networks below $286/175$ by using inequalities that are only valid for global minima of the potential function. Note that the example in Figure 2(a) has a price of stability of 1, since its social optimum is itself a Nash equilibrium.

So far, the best lower bound on the price of stability for three-player games was $74/48 \approx 1.542$ [7]. We can slightly improve this bound by presenting a game with three players whose price of stability is $1769/1126 \approx 1.571$. Consider the network with 5 vertices shown in Figure 2(b). By exhaustive enumeration of all strategy profiles, one can verify that only the strategy profile in which each player i uses the edge (s_i, t_i) is a Nash equilibrium. The social optimum uses all other edges and has a cost of 1126, while the unique Nash equilibrium has a cost of 1769. This establishes the claimed lower bound on the price of stability in undirected Shapley network design games with three players.

5 Conclusions

We gave an upper bound for the price of stability for an arbitrary number of players k in undirected graphs. Our bound is smaller than H_k for every k and tight for two players. For three players, we showed that the upper bound of $286/175 \approx 1.634$ by Bilò and Bove [8] is tight for both the potential-optimal price of stability and anarchy. We also improved the lower bound to $1769/1126 \approx 1.571$ for the price of stability in this case.

Asymptotically, a wide gap remains between the upper bound of the price of stability that is of order $\log k$, and the best known lower bound construction by Bilò et al. [5] that approaches the constant $348/155 \approx 2.245$. It is unclear where the correct answer lies within this gap, in particular since bounds of $O(\log k / \log \log k)$ [9] and $O(\log \log k)$ [6] emerge for restrictions of the problem. Our bound approaches H_k for a growing number of players k . It would already be interesting to know whether the price of stability is asymptotically below H_{k-c} or $H_{k/c}$ for some (or even any) positive constant $c \in \mathbb{N}$.

Acknowledgements. This work has been partially supported by the Swiss National Science Foundation (SNF) under the grant number 200021.143323/1. We thank Julian Mestre for pointing out a mistake in the proof of Lemma 1 in an earlier version of this paper.

References

1. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38, 1602–1623 (2008)
2. Tardos, É., Wexler, T.: Network formation games. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press (2007)
3. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
4. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14, 124–143 (1996)
5. Bilò, V., Caragiannis, I., Fanelli, A., Monaco, G.: Improved lower bounds on the price of stability of undirected network design games. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) *SAGT 2010*. LNCS, vol. 6386, pp. 90–101. Springer, Heidelberg (2010)
6. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the price of stability for designing undirected networks with fair cost allocations. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006*. LNCS, vol. 4051, pp. 608–618. Springer, Heidelberg (2006)
7. Christodoulou, G., Chung, C., Ligett, K., Pyrga, E., van Stee, R.: On the price of stability for undirected network design. In: *Proc. 7th International Workshop on Approximation and Online Algorithms*, pp. 86–97 (2009)
8. Bilò, V., Bove, R.: Bounds on the price of stability of undirected network design games with three players. *Journal of Interconnection Networks* 12, 1–17 (2011)
9. Li, J.: An upper bound on the price of stability for undirected shapley network design games. *Information Processing Letters* 109, 876–878 (2009)
10. Blume, L.E.: The statistical mechanics of best-response strategy revision. *Games and Economic Behavior* 11, 111–145 (1995)
11. Asadpour, A., Saberi, A.: On the inefficiency ratio of stable equilibria in congestion games. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 545–552. Springer, Heidelberg (2009)
12. Kawase, Y., Makino, K.: Nash equilibria with minimum potential in undirected broadcast games. In: Rahman, M. S., Nakano, S.-I. (eds.) *WALCOM 2012*. LNCS, vol. 7157, pp. 217–228. Springer, Heidelberg (2012)