

On the Price of Stability for Undirected Network Design

George Christodoulou¹, Christine Chung^{2,*}, Katrina Ligett³,
Evangelia Pyrga¹, and Rob van Stee^{1,**}

¹ Max-Planck-Institut für Informatik, Saarbrücken, Germany
{gchristo, pyrga, vanstee}@mpi-inf.mpg.de

² Connecticut College, New London, CT, USA
cchung@conncoll.edu

³ Cornell University, Ithaca, NY, USA
katrina@cs.cornell.edu

Abstract. We continue the study of the effects of selfish behavior in the network design problem. We provide new bounds for the price of stability for network design with fair cost allocation for undirected graphs. We consider the most general case, for which the best known upper bound is the Harmonic number H_n , where n is the number of agents, and the best previously known lower bound is $12/7 \approx 1.778$.

We present a nontrivial lower bound of $42/23 \approx 1.8261$. Furthermore, we show that for two players, the price of stability is exactly $4/3$, while for three players it is at least $74/48 \approx 1.542$ and at most 1.65. These are the first improvements on the bound of H_n for general networks. In particular, this demonstrates a separation between the price of stability on undirected graphs and that on directed graphs, where H_n is tight. Previously, such a gap was only known for the cases where all players have a shared source, and for weighted players.

1 Introduction

The effects of selfish behavior in networks is a natural problem with long-standing and wide-spread practical relevance. As such, a wide variety of network design and connection games have received attention in the algorithmic game theory literature (for a survey, see [1]).

One natural question is how much the users' selfish behavior affects the performance of the system. Koutsoupias and Papadimitriou [2,3] addressed this question using a worst-case measure, namely the *Price of Anarchy* (PoA). This notion compares the cost of the worst-case Nash equilibrium to that of the social optimum (the best that could be obtained by central coordination). From an optimistic point of view, Anshelevich et al. [4] proposed the *Price of Stability* (PoS), the ratio of the lowest Nash equilibrium cost to the social cost, as a measure of the minimal effect of selfishness.

There has been substantial work on the PoA for *congestion games*, a broad class of games with interesting properties originally introduced by Rosenthal [5]. Congestion games nicely model situations that arise in selfish routing, resource allocation and

* Some of this research was conducted while this author was visiting the Max-Planck-Institut für Informatik, Saarbrücken, Germany.

** Research supported by German Research Foundation (DFG).

network design problems, and the PoA for these games is now quite well-understood [6,7,8,9]. By comparison, much less work has been done on the PoS: The PoS for network design games has been studied by [4,10,11,12,13], while the PoS for routing games¹ was studied by [4,8,14]. However, PoA techniques cannot easily be transferred to the study of PoS. New techniques thus need to be developed; this work moves toward this direction.

The particular network design problem we address here is the one which was initially studied by Anshelevich et al. [4], sometimes referred to as the fair cost sharing network design (or creation) game. In it, each player has a set of endpoints in a network that he must connect; to achieve this, he must choose a subset of the links in the network to utilize. Each link has a cost associated with it, and if more than one player wishes to utilize the same link, the cost of that link is split evenly among the players. Each player's goal is to pay as little as possible to connect his endpoints. The global social objective is to connect all player's endpoints as cheaply as possible.

Anshelevich et al. [4] showed that if G is a directed graph, the price of anarchy is equal to n , the number of players, whereas the price of stability is exactly the n th harmonic number H_n . The upper bound is proven by using the fact that our network design game, and in fact any congestion game, is a potential game. A *potential game*, first defined by Monderer and Shapley [15], is a game where the change to a player's payoff due to a deviation from a game solution can be reflected in a *potential function*, or a function that maps game states to real numbers.

This upper bound of H_n holds even in the case of undirected graphs (since the potential function of the game does not change when the underlying graph is undirected), however the lower bound does not. Hence the central open question we study is:

What is the price of stability in the fair cost sharing network design game on undirected graphs?

In the case of two players and a single common sink vertex, Anshelevich et al. [4] show that the answer is $4/3$. Some further progress has also more recently been made toward answering this question. Fiat et al. [12] showed that in the case where there is a single common sink vertex and every other vertex is a source vertex, the price of stability is $O(\log \log n)$. They also give an n -player lower bound instance of $12/7$ [16]. For the more general case where the agents share a sink but not every vertex is a source vertex, Li [13] showed an upper bound of $O(\log n / \log \log n)$. Chen and Roughgarden [10] studied the price of stability for the *weighted* variant of the game, where each player pays a fraction of each edge cost proportional to her weight. Albers [11] showed that in this variant, the price of stability is $\Omega(\log W / \log \log W)$, where W is the sum of the players' weights.

Our contributions. We show for the first time that the price of stability in undirected networks is definitively different from the one for directed networks in the general case

¹ Both cost-sharing network design games and network routing games fall in the class of congestion games and they differ only in the edge cost functions. Cost sharing network design games come together with decreasing cost functions on the edges, e.g. $c_e(x) = c_e/x$, while routing games come with increasing latency functions, e.g. $c_e(x) = c_e \cdot x$.

(where all players may have distinct source and destination vertices). In particular, we show that PoS is exactly $4/3$ for two agents (strictly less than PoS in the directed case, which is $H_2 = 3/2$), while for three agents it is at least $74/48 \approx 1.542$ and at most 1.65 (again strictly less than PoS in the directed case, which is $H_3 = 11/6$). Furthermore, we show that the price of stability for general n is at least $42/23 > 1.8261$, improving upon the previous bound due to Fiat et al. [12].

1.1 The Model

We are given an underlying network, $G = (V, E)$, where V is the set of vertices and E is the set of edges in the network. Each player $i = 1 \dots n$ has a set of two nodes (endpoints) $s_i, t_i \in V$ to connect. We refer to s_i as the *source* endpoint of player i and t_i as the *destination* or *sink* endpoint of player i . The strategy set of each player i consists of all sets of edges $S_i \subseteq E$ such that S_i connects all the vertices in T_i . There is a cost c_e associated with each edge $e \in E$. The cost to player i of a solution $S = (S_1, S_2, \dots, S_n)$ is $C_i(S) = \sum_{e \in S_i} c_e/n_e$ where n_e is the number of players in S who chose a strategy that contains e . Each player i wants to minimize $C_i(S)$. The global objective is minimize $\sum_{i=1}^n C_i(S)$.

2 A Lower Bound of 1.826

Consider a 3 by N grid for some large N . There are three nodes and two horizontal edges in every row. The levels are numbered starting from the bottom. We denote the horizontal edges on level i by L_i and R_i (from left to right). The nodes on level i are denoted by v_{ij} ($j = 1, 2, 3$) and the vertical edges connecting level i to level $i + 1$ are denoted by e_{ij} ($j = 1, 2, 3$). Each node v_{ij} for $i = 1, \dots, N - 1$ and $j = 1, 2, 3$ is the source of some agent $p_{i,j}$, who has node $v_{i+1,j}$ as its sink. We say that player $p_{i,j}$ *starts at level i* . Also we will call player $p_{i,j}$ the *owner* of edge $e_{i,j}$, with $p_{i,j}$ *owning* only edge $e_{i,j}$ (one of the possible paths for a player to reach its sink is to use just the edge it owns).

Horizontal edges cost $6 + \varepsilon$ and $5 + \varepsilon$, vertical edges cost 12, 15, and 15 (from left to right), where ε is a small positive number. We do not refer to ε in the calculations, but simply state when relevant that the costs of horizontal edges are “more than” 6 and 5, respectively. One motivation for choosing the numbers as we do is shown in Figure 1, right.

Proof outline. It is possible to connect the sources and sinks of all the players by using all the horizontal edges and only the vertical edges on the left. For large N and small ε , the cost of this tends to 23 per level.

Our goal is to show that in a Nash equilibrium, all players use the direct link between their source and their sink. Let us assume that some players deviate from this. We start by considering a level i which is not visited by any agents with higher sources, and also not by agents that have lower sinks. In Lemma 1, we show that any agent that reaches such a level moves immediately to its sink.

We prove in Lemma 3 that as long as no agent uses any edge below its source vertex, all agents move straight to their sinks. Section 2.3 is devoted to showing that it is indeed




Fig. 1. On the left are two levels in our construction. The situation on the right is not a Nash equilibrium because of the added ϵ 's on the horizontal edges. The numbers in the right figure give the costs for each agent that uses these edges.

the case that no player moves below its source vertex. To do so, we first bound the number of players that can reach a given level from below in Lemma 4. We find that there can be at most two such agents. This in turn helps us to show in Lemmas 5-7 that players that move below their sources would have to pay too much for their paths, thus showing that no agent moves below their starting levels. This immediately gives us our result, which is summarized in Theorem 1.

Due to lack of space, certain proofs are omitted.

Observation 1. *In a Nash equilibrium, all player paths are acyclic, and the graph that is formed by the paths of any pair of players is acyclic as well.*

Thus, whenever we find a cycle of one of these types, we know that this is not a Nash equilibrium.

Observation 2. *If $e_{i,j}$ is used by any player, it is used by player $p_{i,j}$.*

Proof. If this were not true, the path of any player using $e_{i,j}$ together with the path of $p_{i,j}$ forms a cycle.

Definition 1. *We call a node a terminal if it has a single incident edge at the graph induced by all the player paths in a Nash equilibrium.*

Observation 3. *Consider the graph induced by all the player paths in a Nash equilibrium. (This graph is not necessarily acyclic!) Any path which leads to a terminal and where all intermediate nodes have degree 2 is used only by agents with sources and/or sinks on that path. In particular, an edge which leads to a terminal is used by at most two players: the one with its source at the terminal, and the one with its sink at the terminal.*

Observation 4. *Any player that uses a vertical edge $e_{i,j}$ without owning it, must also be using at least one horizontal edge in some level $i' \leq i$, and one in some level $i'' \geq i + 1$.*

Players on the left We begin by making sure that players on the left always use the edge they own (the direct link between their source and sink). To do so, for all levels i , we substitute $e_{i,1}$ by a path of three edges $\hat{e}_{i,1}, \hat{e}_{i,2}, \hat{e}_{i,3}$ each of which has cost 4 (and thus the path of the three edges together has cost 12). Player $p_{i,1}$ is also substituted by three players $\hat{p}_{i,j} (j = 1, 2, 3)$, with $\hat{p}_{i,j}$ having as source and sink the lower and upper endpoints of edge $\hat{e}_{i,j}$, respectively. (Player $\hat{p}_{i,1}$ has node $v_{i,1}$ as its source and player $\hat{p}_{i,3}$ has node $v_{i+1,1}$ as its sink.) One can now see that the players $\hat{p}_{i,j} (j = 1, 2, 3)$ will never deviate from their own edges; each such player would have to share *two* edges of cost 4 with only their owners, since its sink and/or its source would be terminals. Given that these players will never deviate, we will treat them as one player $p_{i,1}$, and the path $\hat{e}_{i,1}, \hat{e}_{i,2}, \hat{e}_{i,3}$ as the single edge $e_{i,1}$, with $p_{i,1}$ using edge $e_{i,1}$ in any Nash Equilibrium.

2.1 Separators

Definition 2. *Level i is called a separator if no player with source above level i and no player with sink below level i visits level i .*

Lemma 1. *Let level i be a separator. Let $p = p_{i-1,2}$ and $p' = p_{i-1,3}$.*

1. *If player p arrives at level i via edge $e_{i-1,1}$ ($e_{i-1,3}$), and there is no other player which uses that edge besides its owner, then p uses edge L_i (R_i), and shares that edge with at most 2 players.*
2. *If player p' arrives at level i via $e_{i-1,1}$ together with p , it uses L_i and R_i , and pays at least 4 for them. In particular, there are at most 4 agents on L_i .*
3. *If p' arrives at level i via edge $e_{i-1,2}$, it uses edge R_i , and pays at least $5/2$ for it.*

Lemma 2. *Let level i be a separator. Let $p = p_{i-1,2}$ and $p' = p_{i-1,3}$. Assume that player p' does not move below its source. If it arrives at level i via edge $e_{i-1,1}$, then there is some other player which uses $e_{i-1,1}$ besides its owner.*

Proof. The first three edges on the path of p' are R_{i-1}, L_{i-1} , and $e_{i-1,1}$ in this case. Consider agent p . It cannot use edge $e_{i-1,3}$ (in that case, by Observation 2, player p' would use it too) or edge $e_{i-1,1}$ (assumption) in this case, so p uses edge $e_{i-1,2}$. This means that p' cannot use edge L_i (Observation 1). It also implies that edges $e_{i-1,2}$ and $e_{i-1,3}$ are used by at most three players, since they are not used by any left player, any player with source at level $i + 1$ or higher, or p , leaving only $p_{i,2}, p_{i,3}$ and $p' = p_{i-1,3}$ as candidates. Therefore, the cost of these edges is at least 5 to any player. Player p' must use one of them. In addition, p' pays 6 for edge $e_{i-1,1}$, and also 6 for edge $e_{i,1}$ as long as player $p_{i,2}$ or $p_{i,3}$ do not join it. But in that case, the total cost of p' is at least $6 + 6 + 5 > 15$, a contradiction. So $p_{i,2}$ or $p_{i,3}$ must be on $e_{i,1}$. Only one of them can in fact be there since one of the vertical edges $e_{i,2}$ and $e_{i,3}$ must be in use. This means that the cost for $e_{i,1}$ is 4 in this case. However, in this case, the edge that p' uses to come back down to level i costs 7.5. We conclude that if p' pays 6 for $e_{i,1}$, its total cost is at least $6 + 6 + 5 > 15$, and otherwise, its total cost is at least $6 + 4 + 7.5 > 15$. In both cases, this implies that this is not a Nash equilibrium.

Lemma 3. *Consider a Nash equilibrium in which no agent uses any edge below its source. Then all agents move straight to their sinks.*

Proof. We twice use induction. We first show that all players on the right move straight to their sinks, while players in the middle either move straight to their sink or move left, up, and immediately right. Using this, we then show that all players in the middle move straight to their sink.

Consider first level 1. By the assumption of this Lemma, level 2 is a separator. If player $p_{1,3}$ uses L_1 , player $p_{1,2}$ must do this as well by Lemma 2. In addition, in this case $p_{1,3}$ uses R_1 as well, but $p_{1,2}$ does not, and neither does any other player. Both $p_{1,3}$ and $p_{1,2}$ then use edge $e_{1,1}$, and then $p_{1,3}$ continues via edge L_2 and R_2 by Lemma 1, Case 2. This fixes its entire path. We can now calculate the cost for this path depending on the first edge on the path of $p_{2,2}$.

If this is L_2 , the cost is more than $5 + 3 + 4 + 1.5 + 2.5 = 16$ ($p_{2,2}$ is not on R_2 in this case, and neither is $p_{1,2}$). If the first (and only) edge is $e_{2,2}$, the cost is more than $5 + 3 + 4 + 2 + 2.5 = 19$. If the first edge is R_2 , the cost is more than $5 + 3 + 4 + 3 + 2.5 = 17.5$ ($p_{2,3}$ is not on R_2 in this case, because $p_{2,2}$ uses $e_{2,3}$, Observation 2). In all cases, this is too much.

This shows that $p_{1,3}$ does not use edge L_1 . Suppose that $p_{1,3}$ uses R_1 . It then uses $e_{1,2}$ together with $p_{1,2}$ (Observation 2). Since R_2 is used by at most one of the players $p_{2,2}$ and $p_{2,3}$ (by Observation 1 and because no agents move down below their source), $p_{1,3}$ pays more than $5 + 15/2 + 5/2 = 15$, a contradiction. The exact same calculation shows that $p_{1,2}$ does not use R_1 .

The only case left open is the one where $p_{1,2}$ uses L_1 , but $p_{1,3}$ does not. However, in this case, due to Lemma 1, Case 1, it also uses L_2 to reach its sink, making level 3 a separator, because level 3 is not visited by $p_{1,2}$ or $p_{1,3}$. Note that if $p_{1,2}$ does move directly to its sink from its source, then level 3 is a separator too.

We can now continue the proof by induction. Consider a level i and assume that all lower players on the right move straight to their source, where lower players in the middle might deviate and use the left edge. Also by induction, assume that level $i + 1$ is a separator, so that we can use the same lemmas as in the base case. Compared to the calculations above for the case where $p_{1,3}$ uses L_1 , the only change is that edge L_i might cost only $2 + \varepsilon/3$ instead of $3 + \varepsilon/2$, since at most one additional agent ($p_{i-1,2}$) may be using it. This still gives a total cost of more than 15 in all cases, completing the first part of the proof.

We can now prove, also using induction, that agents in the middle move straight to their source. If $p_{1,2}$ uses L_1 , it pays more than $6 + 6 + 3 = 15$ since L_1 now costs more than 3, so it does not do that. By induction, if no player below level i deviates, we find the same calculation for any middle player that moves left. This completes the proof.

2.2 The Number of Agents That Visit a Certain Level

Definition 3. Let S_ℓ be a set of players that visit a horizontal edge at or above level ℓ and that all have sinks at or below level ℓ .

Observation 4 implies the following Corollary:

Corollary 1. For any level i , any player with source below i that uses an edge $e_{i-1,j}$, $j \in \{1, 2, 3\}$, without owning it, belongs to S_i .

Lemma 4. *We have $|S_\ell| \leq 2$.*

Corollary 2. *Any horizontal edge is used by at most four agents, any vertical edge by at most five.*

2.3 Agents Do Not Move Down

Due to Lemma 3, all we need to show is that no player moves below its starting level in a Nash equilibrium. Consider the topmost level i such that there is a player A , with source at level i , that moves below i . Denote the other player that has its source on level i and that does not start on the left by A' . (Note that the player with source $v_{i,1}$ never deviates). A must visit levels below i before it reaches level $i + 1$. Otherwise, either i reaches its sink before going down to $i - 1$, or it will have to form a cycle within its path to go back up to $i + 1$. Similarly, since A goes both below and above level i , it cannot use both L_i and R_i . In the following, we will be repeatedly making use of Lemma 4 and Corollary 2, and the fact that no player with source above level i ever visits a level $i' \leq i$ (by definition of A, A').

Lemma 5. *A does not move first horizontally and then down.*

Proof. Assume that A uses first one of the horizontal edges of level i and then immediately goes down. Since A has to go back up to level i , it creates a path connecting all three nodes of level i using only edges incident to nodes of levels $j \leq i$. This implies that there is no player p with source at level $i - 1$ or below, that visits level $i + 1$. To see this, note that after reaching $i + 1$, p would eventually have to go back down to level i , thus creating another path connecting two nodes of i , this time containing only edges incident to nodes in levels $j' \geq i$ (with at least one vertical edge incident to a node of level $i + 1$). Therefore, the paths of A and p would form a cycle. By definition of A , there is also no player with source above level i that visits level i .

Let c_1 be the column that A starts from, $c_2 \neq c_1$ the column it reaches after using the first horizontal edge, and c_3 the remaining column of the grid. Note that since A uses a horizontal edge of level i , one of c_1, c_2 must be the middle column. A cannot create a cycle going from its source back to level i , therefore it must use edge $e = e_{i-1, c_3}$. Moreover, edge $(v_{i, c_3}, v_{i, 2})$ is not used by any player, otherwise a cycle with A 's path would be formed. Therefore, any player on e that does not own it (including A), must also use $e' = e_{i, c_3}$. Given that no player with source below i visits level $i + 1$, and no player with source above i visits level $i - 1$, the edges e', e'' can only be used by the owners and A, A' . Therefore, A pays at least $2 \cdot \frac{12}{3} = 8$ for them.

Consider now the first edge that A uses to reach level $i - 1$. By Corollary 2 there are at most 5 players using it, and thus A pays at least $\frac{12}{5} > 2$ for it. Finally, A visits both the first column and the third column of the grid, therefore it must use at least two ‘‘right’’ horizontal edges (of cost 5), and at least two ‘‘left’’ horizontal edges (of cost 6), each of which can be used by at most four players (by Corollary 2). Thus, A pays at least $2 \cdot \frac{6+5}{4} = 5.5$ for horizontal edges, implying a total cost more than $8+2+5.5 > 15$, a contradiction.

Lemma 6. *If A starts in the middle column, it does not move straight down from its source.*

Lemma 7. *If A starts in the right column, it does not move straight down from its source.*

Proof. Assume that A goes straight down from its source. We denote by e that first edge down (i.e., $e = e_{i-1,3}$). Let e' be the edge that A uses to reach level i again, after going down. By Lemma 5 and Lemma 6, A' does not move down which means that A' does not use e, e' . Any other player using them, apart from A and the owners, will belong to S_i (remember that no player with source above i visits a level below $i + 1$). Therefore A shares e, e' with at most 3 more players (the owners and two more players that will belong to S_i). Let e'' be the edge A uses to reach level $i + 1$ from i . Any player on e'' (apart from the owner) will belong to S_{i+1} together with A . Again, since $|S_{i+1}| \leq 2$, A shares e'' with at most two more players (the owner and one more player that will belong to S_{i+1}). If any of e', e'' is in the left column, then the path of A must cross from the right side of the grid to the left and back, implying a total cost of at least $15/4$ (for e) $+12/4$ (for e') $+12/3$ (for e'') $+2 \cdot \frac{6+5}{4} > 15$. If, on the other hand, none of e', e'' are in the left column, then the total cost of A is more than $15/4$ (for e) $+15/4$ (for e') $+15/3$ (for e'') $+2 \cdot 5/4 = 15$.

Theorem 1. *The price of stability in undirected networks is at least $42/23 > 1.826$.*

Proof. Due to Lemma 5, Lemma 6 and Lemma 7, no agents move down below their source. Therefore, by Lemma 3, all agents move straight to their sink in the (unique) Nash equilibrium. On every level, the total cost of the agents in the Nash equilibrium is $12 + 15 + 15 = 42$, whereas the optimal cost is only $12 + 6 + 5 = 23$. The optimal solution has an additional cost of 11 for the two horizontal edges on level 1, but this cost is negligible for large N .

3 Two and Three Players

We will describe here a lower and an upper bound for three players, as well as an unconditional upper bound for two players. Again, some proofs are omitted due to lack of space.

Lower bound for three players. Figure 2 shows a three-player instance where the best Nash equilibrium has cost $37/24$ times that of OPT. Node s_i, t_i is the source, destination, respectively of player $i, i \in \{1, 2, 3\}$. The optimal solution would only use the edges $(s_1, s_2), (s_2, s_3), (s_3, t_1), (t_1, t_2)$, while the Nash solution uses the direct edges $(s_1, t_1), (s_2, t_2), (s_3, t_3)$. The cost of the optimal solution sums then up to $48 + 4\epsilon$, while the Nash Equilibrium solution has cost 74. We have therefore the following theorem.

Theorem 2. *In the fair cost sharing network design game with three players, the price of stability is at least $74/48 \approx 1.5417$.*

Upper bound for three players. Given an instance of our problem, let OPT refer to an optimal solution. We refer to the union of the players' paths at OPT as the *OPT*




Fig. 2. A three-player instance with price of stability more than 1.54

graph. Recall that our game is a potential game, with potential function $\Phi(X) = \sum_{e \in E} c_e H(X_e)$ where c_e is the cost of edge e , $H(x)$ is the x th harmonic number, X is a game state or solution, and X_e is the number of players on edge e in X . Let N be a potential minimizing Nash solution (or, alternatively, N can be defined as a Nash solution reached by starting from OPT and making alternating best-response moves). Hence, we have

$$\Phi(N) \leq \Phi(OPT). \tag{1}$$

We now give names for various sets of edges, each of which may or may not be empty. Let $A, B,$ and C be the sets of edges that player 1, player 2, and player 3 (respectively) use alone in N . Let S_{ij} for $i = 1 \dots 2$ and $j = i + 1 \dots 3$ be the set of edges that players i and j alone share in N . Let S_{123} be the set of edges that all three players share in N . Let $A^*, B^*, C^*, S_{12}^*, S_{13}^*, S_{23}^*$ and S_{123}^* be defined analogously for OPT . We will also use the same names to refer to the total cost of the edges in each set.

Let $C(X)$ refer to the cost of the solution X and let $C_i(X)$ refer to the cost just to player i of the solution X . By definition, we have

$$\begin{aligned} C(N) &= A + B + C + S_{12} + S_{23} + S_{13} + S_{123} \\ C(OPT) &= A^* + B^* + C^* + S_{12}^* + S_{23}^* + S_{13}^* + S_{123}^* \\ C_1(N) &= A + \frac{S_{12}}{2} + \frac{S_{13}}{2} + \frac{S_{123}}{3} \\ C_2(N) &= B + \frac{S_{12}}{2} + \frac{S_{23}}{2} + \frac{S_{123}}{3} \\ C_3(N) &= C + \frac{S_{13}}{2} + \frac{S_{23}}{2} + \frac{S_{123}}{3} \end{aligned}$$

Lemmas 8, 9 show how to bound the POS depending on whether $S_{123}^* > 0$ or not.

Lemma 8. *In the fair cost sharing network design game with three players, if all three players share at least one edge of positive cost in the optimal solution, the price of stability is at most $33/20 = 1.65$.*

Proof. First observe that the edges in the set S_{123}^* must form a contiguous path, that is, once the three players' paths in the OPT graph merge, as soon as one player's path

breaks off, the three may never merge again. (Otherwise the OPT graph would have a cycle, contradicting the fact that it is an optimal solution.) Without loss of generality, we can exchange the labels on the endpoint vertices so that the three endpoints on the same side of the edges in S_{123}^* are all source endpoints, and the three endpoints on the other side are all destination endpoints.

Then observe that at least one of S_{12}^* , S_{23}^* , and S_{13}^* must be empty. Otherwise the OPT graph would have a cycle, contradicting the definition of OPT. Without loss of generality, we assume that S_{13}^* is empty, hence $S_{13}^* = 0$ and $C(OPT) = A^* + B^* + C^* + S_{12}^* + S_{23}^* + S_{123}^*$.

We know by definition of N that each player i pays not more at N than by unilaterally defecting to any alternate $s_i - t_i$ connection path. The right hand sides of each of the following inequalities represents an upper bound on the cost of a feasible alternate $s_i - t_i$ path for each player i . The existence of these alternate paths depends on the assumption that the OPT graph is connected and $S_{13}^* = 0$.

$$C_1(N) \leq A^* + B^* + S_{23}^* + \frac{B}{2} + \frac{S_{12}}{2} + \frac{S_{23}}{3} + \frac{S_{123}}{3} \quad (2)$$

$$C_2(N) \leq B^* + A^* + S_{23}^* + \frac{A}{2} + \frac{S_{12}}{2} + \frac{S_{13}}{3} + \frac{S_{123}}{3} \quad (3)$$

$$C_2(N) \leq B^* + C^* + S_{12}^* + \frac{C}{2} + \frac{S_{23}}{2} + \frac{S_{13}}{3} + \frac{S_{123}}{3} \quad (4)$$

$$C_3(N) \leq C^* + B^* + S_{12}^* + \frac{B}{2} + \frac{S_{23}}{2} + \frac{S_{12}}{3} + \frac{S_{123}}{3} \quad (5)$$

To interpret the above inequalities intuitively, consider for example the first inequality. It states the fact that player 1 pays an amount at Nash that is at most the cost of unilaterally deviating and instead taking the path in the OPT graph from s_1 to s_2 where player 2's OPT path begins (possibly using edges from A^* , B^* , and S_{23}^*), then following along player 2's path in N from s_2 to t_2 (using edges from B , S_{12} , S_{23} , and S_{123}), then taking edges in the OPT graph from t_2 to t_1 (again possibly using edges from A^* , B^* , and S_{23}^*). The costs of S_{12}^* and S_{123}^* need not be included in the right-hand side of the first inequality for the following reasoning. Recall that by assumption, source vertices are on one side of the edges in S_{123}^* and sink vertices are on the other side of the edges in S_{123}^* , so traversing any edges in S_{123}^* is not necessary for player 1 to go from s_1 to s_2 or from t_2 to t_1 in the OPT graph. Also note that the edges in S_{12}^* must be adjacent to the contiguous path formed by edges in S_{123}^* (since otherwise, the OPT graph would contain a cycle), and so in fact, s_1 and s_2 are on one side of $S_{12}^* \cup S_{123}^*$, while t_1 and t_2 are on the other.

From inequality (1) and the assumption that $S_{13}^* = 0$, we can say

$$A + B + C + \frac{3}{2}(S_{12} + S_{13} + S_{23}) + \frac{11}{6}S_{123} \leq A^* + B^* + C^* + \frac{3}{2}(S_{12}^* + S_{23}^*) + \frac{11}{6}S_{123}^*. \quad (6)$$

Scaling the inequalities 2 and 5 each by $10/99$, 3 and 4 each by $8/99$, and 6 by $6/11$, then summing all five resulting inequalities yields




Fig. 3. A sample OPT graph. Each edge is labeled with the name of the set of edges it belongs to. Each edge here may represent a sequence of edges forming a path. Note that more generally, any of the sets $A^*, B^*, C^*, S^*_{12}, S^*_{23}$, and S^*_{13} could be empty.

$$\begin{aligned} \frac{20}{33}(A + B + C) + \frac{257}{297}S_{13} + \frac{245}{297}(S_{12} + S_{23}) + S_{123} \\ \leq \frac{8}{11}(A^* + C^*) + \frac{10}{11}B^* + S^*_{12} + S^*_{23} + S^*_{123}. \end{aligned} \tag{7}$$

Hence $20/33C(N) \leq C(OPT)$.

Lemma 9. *In the fair cost sharing network design game with three players, if no positive-cost edge is shared by all three players in the optimal solution, the price of stability is at most 3/2.*

We are now ready to present our main theorem of this section.

Theorem 3. *In the fair cost sharing network design game with three players, the price of stability is at most $33/20 = 1.65$.*

Proof. All possible OPT graph structures are handled by Lemmas 9 and 8. The worst upper bound for price of stability over these two exhaustive cases is that given by Lemma 8.

Upper bound for two players. Anshelevich et al. [4] gave a two player lower bound instance for our problem showing that the price of stability is at least $4/3$. They then show that if both players share a sink, the price of stability is at most $4/3$. The following theorem states an unconditional two-player upper bound on the price of stability of $4/3$.

Theorem 4. *In the fair cost sharing network design game with two players, price of stability is at most $4/3$.*

4 Conclusions

The lower bound instance that we use for large n could be generalized by adding more columns. However, it seems that this would require a significantly longer and more involved proof. More importantly, we believe that even with an unbounded number of columns we could only show a lower bound of a small constant. Hence, the question of whether the price of stability grows with n remains open. We conjecture that it is in fact constant.

Acknowledgments. The authors would like to thank Susanne Albers, Elliot Anshelevich, Jian Li, Kirk Pruhs and Alexander Souza for interesting discussions.

References

1. Tardos, E., Wexler, T.: Network Formation Games and the Potential Function Method. In: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
2. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
3. Papadimitriou, C.H.: Algorithms, games, and the internet. In: Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC), pp. 749–753 (2001)
4. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: 45th Symposium on Foundations of Computer Science (FOCS), pp. 295–304 (2004)
5. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
6. Roughgarden, T., Tardos, E.: How bad is selfish routing? *J. ACM* 49(2), 236–259 (2002)
7. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), pp. 67–73 (2005)
8. Christodoulou, G., Koutsoupias, E.: On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)
9. Awerbuch, B., Azar, Y., Epstein, A.: Large the price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), pp. 57–66 (2005)
10. Chen, H.L., Roughgarden, T.: Network design with weighted players. In: SPAA 2006: Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures, pp. 29–38. ACM, New York (2006)
11. Albers, S.: On the value of coordination in network design. In: SODA 2008: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 294–303. Society for Industrial and Applied Mathematics, Philadelphia (2008)
12. 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)
13. Li, J.: An $O(\log \log \log n)$ upper bound on the price of stability for undirected Shapley network design games. *Information Processing Letters* (2009)
14. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
15. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14, 124–143 (1996)
16. Olonetsky, S.: On the price of stability for designing undirected networks with fair cost allocations. Master’s thesis, Tel-Aviv University (2006)