

On the Price of Stability for Undirected Network Design

George Christodoulou · Christine Chung ·
Katrina Ligett · Evangelia Pyrga · Rob van
Stee

Received: date / Accepted: date

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.

Keywords Algorithmic Game Theory · Price of Stability · Network Design

A preliminary version of this work appeared in [8].

G. Christodoulou
University of Liverpool, Liverpool, UK,
E-mail: gchristo@liv.ac.uk

C. Chung
Connecticut College, New London, CT, USA,
E-mail: cchung@conncoll.edu

K. Ligett
Cornell University, Ithaca, NY, USA,
E-mail: katrina@cs.cornell.edu

E. Pyrga · R. van Stee
Max-Planck-Institut für Informatik, Saarbrücken, Germany, E-mail: {pyrga,vansteel}@mpi-inf.mpg.de.

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 [17]).

One natural question is how much the users' selfish behavior affects the performance of the system. Koutsoupias and Papadimitriou [10, 14] 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. [2] 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 [15]. Congestion games nicely model situations that arise in selfish routing, resource allocation and network design problems, and the PoA for these games is now quite well-understood [16, 7, 6, 3]. By comparison, much less work has been done on the PoS: The PoS for network design games has been studied by [2, 5, 1, 9, 11], while the PoS for routing games¹ was studied by [2, 6, 4]. 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. [2], 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. [2] 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 [12], 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:

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. $c_e(x) = c_e \cdot x$.

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. [2] show that the answer is $4/3$. Some further progress has also more recently been made toward answering this question. Fiat et al. [9] 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 $12/7$ [13]. For the more general case where the agents share a sink but not every vertex is a source vertex, Li [11] showed an upper bound of $O(\log n / \log \log n)$. Chen and Roughgarden [5] 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 [1] 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 (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. [9].

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*




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.

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 $6 + \epsilon$ and $5 + \epsilon$, 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 Our goal is to show that in a Nash equilibrium all players use the direct link between their source and their sink. To do this, we will upper bound the number of players that can be on any horizontal edge. We will tighten this bound gradually, showing in the end that no player can use any horizontal edge in a Nash Equilibrium, and thus prove the claim.

A few useful observations follow.

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

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

Proof If this were not true, then the path of any player using e_{ij} together with the path of $p_{i,j}$ would form a cycle. \square

Definition 1 We call a node a *terminal* if it has degree 1 in the graph induced by the union of 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 (sub)path that leads to a terminal and such that all intermediate nodes have degree 2 is used only by players 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 whose source is the terminal, and the one whose sink is the terminal.

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

Proof Trivial, since otherwise the player's path would be contained in a single column of the grid. This, however, can only happen if the player uses the direct edge that it owns and no other edge. \square

Observation 5 A player with source at level i uses only one edge $e_{i,j}, j \in \{1, 2, 3\}$.

Proof The player's path begins at level i and ends at level $i+1$, therefore it needs to use an odd number of edges $e_{i,j}$. In order to use three it also needs to use the edge it owns, in which case it would use no other edge. \square

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 The Proof

We will now define two sets of players per level. One can see that the second definition is the symmetric version of the first one.

Definition 2 We define the set S_ℓ as the set of all players whose *sink* lies at some level $k \leq \ell$, and who use some horizontal edge of some level $k' \geq \ell$.

Definition 3 We define the set T_ℓ as the set of all players whose *source* lies at some level $k \geq \ell$, and who use some horizontal edge of some level $k' \leq \ell$.

Observation 4 implies the following Corollary:

Corollary 1 Consider some level i and some player p with source at level i' that uses an edge $e_{i,j}$, (for some $j \in \{1, 2, 3\}$) without owning it. If $i' \leq i$, then $p \in S_{i+1}$, while if $i' \geq i$, then $p \in T_i$. In particular, if $i' = i$ then p belongs to $S_{i+1} \cap T_i$.

We will now fix some set S_ℓ and try to bound its size.

Lemma 1 *Let i be a level that contains the source of some player in S_ℓ . Assume that i is not the lowest such level.*

- i. *There is a single player in S_ℓ with source on level i .*
- ii. *Any player that uses an edge $e_{i,j}$ and does not own it is in S_ℓ .*

Proof i. i is not the lowest level with a player in S_ℓ . Therefore, there is some player $q \in S_\ell$ who must use some edge $e_{i,j}$ to reach level $i+1$ and some other edge $e_{i,j'}, j' \neq j$ to go back down to level i on the way to its sink. This means that $p_{i,j}, p_{i,j'} \notin S_\ell$, since, by Observation 2 they only use the edges they own and no horizontal edge. Consequently the statement holds.

- ii. Suppose there is a player $p \notin S_\ell$ who uses $e_{i,j}$ but does not own it. The source of p is not on level i ; as discussed in i, the three players with sources on i either use the edges they own, or belong to S_ℓ . Let again $e_{i,j}, e_{i,j'}$ be the two vertical edges of level i that are used, and let q be some player in S_ℓ with source below i .

Case 1: Assume that the source of p is below level i . p must then reach up to level $i+1$, using one of $e_{i,j}, e_{i,j'}$, and later return again to i using the other in order to reach its sink. Notice that these two edges will also be used by player $q \in S_\ell$. q reaches up to level ℓ while p does not (otherwise it would also use some horizontal edge at or above ℓ , and thus belong to S_ℓ); this, however, forms a cycle using the paths of p, q above nodes $v_{i+1,j}, v_{i+1,j'}$, a contradiction.

Case 2: Assume that the source of p is above level i . Then also its sink is above i , meaning that it must return back to $i+1$ at some point after reaching i . Thus, again, it must use both $e_{i,j}, e_{i,j'}$. But in that case, we again find a cycle if we combine the paths of p and q (p forms a continuous path linking nodes $v_{i,j}, v_{i,j'}$ from below, while q forms a continuous path linking those two nodes from above).

In both cases we reach a contradiction and the statement holds. \square

Lemma 2 $|S_\ell| \leq 3$. *If the lowest level which contains a source of a player in S_ℓ only contains one such source, then $|S_\ell| \leq 2$.*

Proof Let i be the lowest level that contains a source of a player in S_ℓ .

Assume first there is a unique player $p \in S_\ell$ whose source is on i . Let L_p be the set of levels in the path of p that contain sources of other players that also belong to S_ℓ . p must traverse two edges $e_{k,j}$ for each $k \in L_p$, and by Lemma 1, only players in S_ℓ (and the corresponding owners) traverse them. The cost of p for these edges is then at least $(12+15)\left(\frac{1}{3} + \frac{1}{4} + \dots + \frac{1}{|S_\ell|+1}\right) > 15$ if $|S_\ell| > 2$.

Assume now that level i contains the sources of two players that belong to S_ℓ , say p, p' , and assume that $|S_\ell| > 3$ (and hence $i < \ell - 1$). Of course, none of $e_{i,2}, e_{i,3}$ is used, while $e_{i,1}$ is only used by p, p' and its owner: any other player with source below (above) i would have to use another edge $e_{i,j}, j \in \{2, 3\}$ to return to level i ($i+1$). Therefore, p, p' each pay $12/3 = 4$ for using $e_{i,1}$. Since $\ell > i+1$, they must continue going upwards until they reach level ℓ and can then return to their sinks. But then, similarly to the previous case, if $|S_\ell| > 3$ they pay in total $4 + (12+15)\left(\frac{1}{4} + \frac{1}{5} + \dots + \frac{1}{|S_\ell|+1}\right) > 15$. \square

The symmetric version of Lemma 2, gives the corresponding version that bounds the number of players in T_ℓ , for any $\ell > 0$.

Lemma 3 $|T_\ell| \leq 3$. *If the highest level which contains a source of a player in T_ℓ only contains one such source, then $|T_\ell| \leq 2$.*

Combining Lemmata 2, 3, we obtain the following bounds on the number of players on any edge:

Corollary 2 *There are at most 6 players on any horizontal edge, and at most 7 on any vertical edge.*

Proof Consider some horizontal edge at some level ℓ . Players that use it either have sources below ℓ and thus belong to the set S_ℓ , or they have sources at or above ℓ , in which case they belong to T_ℓ . Lemmata 2, 3 bound the cardinalities of either of the sets to at most 3; thus, any horizontal edge of ℓ can be used at most by 6 players. Similarly, for any $e_{\ell,j}, j \in \{1, 2, 3\}$, it is enough to note that any player that uses it, apart from the owner, belongs to at least one of $S_{\ell+1}, T_\ell$ (Corollary 1). \square

Now, we can show the following.

Lemma 4 S_ℓ *does not contain two players whose sources lie at the same level. Therefore, $|S_\ell| \leq 2$.*

Proof Consider the smallest $\ell > 0$ such that $|S_\ell| = 3$. Again, let i be the lowest level that contains sources of players in S_ℓ . Level i contains the sources of two players in S_ℓ and hence $i < \ell - 1$. Following the lines of the proof of Lemma 2, let p, p' be these two players. More specifically, let p be player $p_{i,2}$, and p' be player $p_{i,3}$.

Since two players from level i do not use their own edges, they will both use edge $e_{i,1}$ to reach level $i + 1$. They cannot use edge L_{i+1} since then p would reach its sink and would not reach up to level $\ell > i + 1$. Moreover, L_{i+1} is not used by any player, since that would create a cycle with p 's path. Therefore, p, p' continue on $e_{i+1,1}$, while they share $e_{i,1}, e_{i+1,1}$ only with the corresponding owners. No other player can use them. Note that only one vertical edge from level i is used. Any player with source above i using $e_{i+1,1}$ and necessarily also $e_{i,1}$ would have no way of returning back to levels above i without creating a cycle; similarly, a player with source below i would have no way of returning back to level i . Therefore, p, p' pay $2 \cdot 12/3 = 8$ for these two edges.

At some point after using $e_{i+1,1}$ they also use some other vertical edge to return to level $i + 1$. That edge cannot be $e_{i+1,1}$, therefore they pay at least $15/7$ for it (Corollary 2). Also, after using $e_{i+1,1}$, p' must return to the rightmost column of the grid, where its source is. Therefore it must use two more horizontal edges (see also Figure 2). Corollary 2 then implies that it will have to pay more than $(6 + 5)/6$ for them.

In total, p' pays at least $8 + 15/7 + 11/6 > 11.9$ just to reach its sink after reaching node v_{i1} .

We can now observe that p' cannot include nodes of level $i - 1$ in its path: if it did, it would have to use two more vertical edges (one to reach level $i - 1$ and another




Fig. 2 Depicting the path of player p' for the proof of Lemma 4 along with lower bounds on the cost of each part. The part from the source $v_{i,3}$ of p' until node $v_{i+2,1}$ is fixed.

to return back to i), at most one of which can be of cost 12. Then its total cost would be more than $11.9 + (12 + 15)/7 > 15$.

p' uses therefore R_i, L_i while p can only use L_i . Any other player on either of these edges must belong to S_i . However, ℓ is the lowest possible level such that $|S_\ell| = 3$; thus $|S_i| \leq 2$. p' pays then more than $5/3$ and $6/4$ for R_i, L_i , respectively, and its total cost becomes more than $11.9 + 5/3 + 6/4 > 15$. This is too expensive, so p' would have used its direct edge instead. Therefore $|S_\ell| \leq 2$ and no two players with sources at the same level can belong to S_ℓ . \square

Again the symmetric version bounds the cardinality of T_ℓ .

Lemma 5 T_ℓ does not contain two players with sources at the same level. Therefore, $|T_\ell| \leq 2$.

Combining now Lemmata 4, 5, we can improve the bound given by Corollary 2 for the number of players that can be on any edge. We therefore have the following:

Corollary 3 Any horizontal (vertical) edge is used at most by four (five) players.

Lemma 6 The path of any player spans at most 3 levels.

Proof For the sake of contradiction, consider some player p whose path spans at least 4 levels. Let i be the level of its source. The source and the sink of p are only one level apart. Therefore, apart from levels $i, i + 1$, any other level that p reaches implies that it uses two additional vertical edges: one to reach that level, and one to return. Also, at most one of these two edges can be of cost 12. Therefore, if it visits at least $k \geq 4$ levels then it must use at least $2(k - 2) + 1$ vertical edges in total (with the “+1” referring to the edge it will use to reach level $i + 1$ from i). Each of these edges can be used by at most 5 players in total, according to Corollary 3. Therefore, its cost for vertical edges only is at least $\frac{12(k-2+1)+15(k-2)}{5}$. It must also use at least two horizontal edges (otherwise it would have to use its own edge). Again, Corollary 3 implies that its cost for these edges will be at least $2 \cdot \frac{5}{4}$. Now assuming that $k \geq 4$, its total cost is at least $\frac{3 \cdot 12 + 2 \cdot 15}{5} + 2.5 > 15$, implying that p would have used its direct edge instead. Therefore $k \leq 3$. \square

This immediately gives the following Corollary.

Corollary 4 *Let p be any player whose source is at level i . p never reaches level $i - 2$, or level $i + 3$.*

Lemma 7 *There is no level ℓ such that $|S_\ell| > 0$.*

Proof Assume that the statement is false. Consider the smallest ℓ such that $S_\ell \neq \emptyset$. Let i be the lowermost level containing the source of some player p , such that $p \in S_\ell$. The other two players of level i use only the edges they own: the leftmost player always does so by construction; if the other player used an edge other than the one it owns, then it would also have to use a horizontal edge in some level $i \geq i + 1$. However the definition of ℓ, p together with Lemma 4 imply that this is not possible.

All players that have source at any level $i' < i$ use the edges they own: Corollary 1 implies that a player with source at i' that does not use the edge it owns would belong to $S_{i'+1}$, with $i' + 1 \leq i < \ell$. But $S_{i'} = \emptyset$ for any $i' < \ell$ due to the definition of ℓ .

Consider now again player p . Let $e_{i,j}$ be the vertical edge of level i it is using. Corollary 1 implies that $p \in T_i$. Any other player p' that uses edge $e_{i,j}$ without owning it will belong to T_i as well (remember that no player with source at any level $i' \leq i$ apart from p uses a horizontal edge). Lemma 5 implies that there can be at most one such player p' . The source of p' is higher than level i , meaning that after it reaches i it needs to return back up. Therefore it uses two of the vertical edges of level i and apart from the corresponding owners, it shares only one of those with p , while only one can cost 12, see also Figure 3 (left). This already costs at least $12/2 + 15/3 = 11$. Of course p' must use also one edge connecting the level of its source with that of its sink, which will cost at least $12/5$ (Corollary 3). Finally, it needs to use at least one horizontal edge of level i or below which can be shared with at most one more player (i.e., p), at cost at least $5/2$. In total, this sums up to $11 + 2.4 + 2.5 > 15$. Therefore, no player is in T_i together with p , and p shares $e_{i,j}$ only with its owner.

This implies that p does not reach level $i - 1$. If it did, it would have to use at least two more vertical edges (first to reach level $i - 1$ and then to leave it again), but given that it is the only player in T_i , it would only share the costs with their owners. That would immediately imply cost more than $2 \cdot 12/2 + 15/2 > 15$. Similarly, p does not reach level $i + 2$; its cost this time would be more than $12/2 + (12 + 15)/5 + 5$ (for a horizontal edge of level i) > 15 .

Hence, p only uses horizontal edges of levels $i, i + 1$, and a single vertical edge $e_{i,j}$, for some $j \in \{1, 2, 3\}$. Also, $\ell = i + 1$. Let us consider some horizontal edge of level $i + 1$ that p uses. Any other player on it will belong either to S_{i+1} , or to T_{i+1} . S_{i+1} consists, however, only of p . Note also that if there is only one player in T_{i+1} then p would have used its direct edge instead: If p uses L_1 , then its cost would be more than $6 + 12/2 + 6/2 = 15$. If p uses R_i (and not L_i), then the vertical edge p uses is of total cost 15, implying it would pay more than $5 + 15/2 + 5/2 = 15$. Therefore, there must be two players in T_{i+1} . Lemma 5 and Corollary 4 imply that one of these players must have source at level $i + 1$, and the other at $i + 2$. Let q be the player with source at level $i + 1$, and q' the one with source at $i + 2$, see also Figure 3 (right).

q only uses one vertical edge of level $i + 1$, while q' must use two. One of these two edges will be shared with q . The other one will only be shared with the owner:




Fig. 3 Possible paths of players in the proof of Lemma 7. Left: A possible path for player p' along with lower bounds on the cost of (some of) the edges it must use. The solid lines indicate edges that would be used both by p and p' . The dashed (dotted) line completes the path of p (p'), along with the solid line. Right: Possible paths of p, q, q' . The solid, dotted, dashed lines correspond to the path of player p, q, q' , respectively.

any other player that uses it without being the owner must belong to T_{i+1} which only consists of q, q' . Of course q' also needs to use some vertical edge connecting level $i+2$ to level $i+3$ (at cost at least $12/5$); it also needs to use at least one horizontal edge of level $i+1$ (possibly) sharing only with p, q (at cost more than $5/3$), as well as some horizontal edge at level $i+3$ or above (at cost more than $5/4$), see also Figure 3 (right). Therefore, its total cost is more than $12/2 + 15/3 + 12/5 + 5/3 + 5/4 > 15$. Hence, q' would have used its direct edge.

This implies that T_{i+1} consists of a single player, causing also p to use its direct edge, and therefore $S_\ell = \emptyset$. \square

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

Proof If a player with source at level i does not use the edge it owns, then $S_{i+1} \neq \emptyset$. Lemma 7 states that this is not possible. Therefore, there is a unique Nash Equilibrium, in which every player uses its own edge.

On every level, the total cost of the players 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 K . \square

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.

Lower bound for three players Figure 4 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.




Fig. 4 A three-player instance with price of stability more than 1.54.

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

Proof Let p_1, p_2, p_3 be the three players (with p_i having to connect s_i to t_i). It is clear that a solution of value $48 + 4\epsilon$ exists. We will show that there is no other Nash Equilibrium besides the one mentioned above of cost 74, i.e., every player p_i uses edge (s_i, t_i) . Note first that edge (s_2, s_3) cannot be used by both players p_2 and p_3 (since their paths would create a cycle, given that they both have to reach t_2).

First, we will show that player p_1 must use (only) the edge (s_1, t_1) . Assume that p_1 does not use the direct edge (s_1, t_1) (and also no other player is using it, as this would create a cycle with p_1).

- Assume first that p_1 is using edge (s_1, s_2) . p_1 must then use either edge (s_2, t_2) or (s_2, s_3) .

p_1 uses (s_2, t_2) Then it must also use (t_1, t_3) . p_2 will also then be on (s_2, t_2) (and will not be using any other edge). If p_3 is on (s_2, t_2) as well, then p_1 must be alone on (s_1, s_2) and (t_1, t_2) , implying a total cost for p_1 equal to $8 + \epsilon + \frac{24}{3} + 8 + \epsilon = 24 + 2\epsilon > 24$ so p_1 would have preferred to use the direct edge (s_1, t_1) instead. If p_3 is not on (s_2, t_2) , then it is also not on (s_1, s_2) (otherwise there would be a cycle with p_2). p_1 's cost would then be $8 + \epsilon + \frac{24}{2} + \frac{8+\epsilon}{2} > 24$. So p_1 again would prefer (s_1, t_1) .

p_1 uses (s_2, s_3) No player uses (s_1, s_3) (since that would create a cycle with p_1), and since (s_1, t_1) is also not in use, p_1 is alone on (s_1, s_2) . Moreover, at most one

of p_2, p_3 can be on (s_2, s_3) . Then p_1 pays at least $8 + \varepsilon + \frac{14+\varepsilon}{2} > 15 + \varepsilon$ in order to reach s_3 . Therefore it would have used edge (s_1, t_1) instead.

Therefore p_1 is not on (s_1, t_1) .

- Suppose p_1 is on (s_1, s_3) . Then it has to use either (s_2, s_3) or (s_3, t_1) .

p_1 uses (s_2, s_3) Then it must also use (s_2, t_2) (implying that p_2 is only using (s_2, t_2)) and (t_1, t_2) . Even if player p_3 was on all edges p_1 uses, its total cost would still be at least $\frac{15+\varepsilon+14+\varepsilon+8\varepsilon}{2} + \frac{24}{3} > 24$, so p_1 would have preferred edge (s_1, t_1) instead.

p_1 uses (s_3, t_1) Consider then player p_2 . Assume that p_2 is not using the direct edge (s_2, t_2) . Given that (s_1, t_1) is not used, p_2 has only two options: Either it uses (s_1, s_2) and (s_1, s_3) , or it uses (s_2, s_3) .

If p_2 is on (s_1, s_2) and (s_1, s_3) , player p_3 cannot have used (s_1, s_2) without creating a cycle either in its own path or with p_2 (remember that edge (s_1, t_1) is not in use). Therefore the cost of p_2 would be at least $8 + \varepsilon + \frac{15+\varepsilon}{2} + \frac{18+\varepsilon}{3} + \frac{8+\varepsilon}{2} > 24$, implying that p_2 would have used the direct edge (s_2, t_2) instead.

If p_2 is on (s_2, s_3) then p_3 cannot be using it. Therefore, p_2 pays at least $14 + \varepsilon + \frac{18+\varepsilon}{3} + \frac{8+\varepsilon}{2} > 24$.

Thus, p_2 will be using (s_2, t_2) .

Given now that p_2 is only using (s_2, t_2) and the fact that p_3 cannot be both on (s_1, s_3) and (s_3, t_1) , the cost of p is at least $15 + \varepsilon + \frac{18+\varepsilon}{2} > 24$ (if p_3 is not on (s_1, s_3)), or $\frac{15+\varepsilon}{2} + 18 + \varepsilon > 24$ (if p_3 is not on (s_3, t_1)).

In all cases, p_1 would therefore prefer to use the direct edge (s_1, t_1) .

We next consider player p_2 . Assume that p_2 is not using the direct edge (s_2, t_2) (and thus p_3 cannot use it either). p_2 will not use (s_2, s_3) , since its cost will then be at least $14 + \varepsilon + \frac{26}{2} > 24$.

Therefore, p_2 uses edge (s_1, s_2) and afterwards it either uses (s_1, s_3) or (s_1, t_1) .

- Assume p_2 uses (s_1, s_3) . In this case p_3 cannot be in either of (s_1, s_3) or (s_1, s_2) , as this would create a cycle (either in its own path, or together with p_2). p_2 would then have to pay at least $8 + \varepsilon + 15 + \varepsilon + \frac{26}{2} > 24$.
- Assume p_2 uses (s_1, t_1) . Consider player p_3 . Assume that p_3 is not using the direct edge (s_3, t_2) , or (s_3, t_1) and then (t_1, t_2) .

Since (s_2, t_2) is not used by any player, p_3 must be using (s_1, t_1) with direction from s_1 to t_1 (just as p_2 does). The cheapest way that p_3 has to reach node s_1 though is via edge (s_1, s_3) . Therefore p_3 would pay in total at least $15 + \varepsilon + \frac{24}{3} + \frac{8+\varepsilon}{2} > 26$, so it would rather use the direct edge (s_3, t_2) instead. Therefore p_3 is either on (s_3, t_2) , or (s_3, t_1) and (t_1, t_2) . As a result, the cost of p_2 is at least $8 + \varepsilon + \frac{24}{2} + \frac{8+\varepsilon}{2} > 24$.

Thus, also p_2 uses the direct edge (s_2, t_2) . Now player p_3 would not use edge (s_2, t_2) since it would require a total cost of at least $14 + \varepsilon + \frac{24}{2} > 26$. It cannot then reach node s_2 as this would create a cycle with p_2 . If it uses (s_1, s_3) it must also use (s_1, t_1) , and pay at least $15 + \varepsilon + \frac{24}{2} + 8 > 26$. Edge (s_3, t_2) results in a lower cost than using both (s_3, t_1) and (t_1, t_2) , and thus p_3 also using the direct edge (s_3, t_2) .

The above imply that the Nash Equilibrium is unique. \square

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 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

$$\Phi(N) \leq \Phi(OPT). \quad (1)$$

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.




Fig. 5 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.

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

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

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

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. \square

Upper bound for two players Anshelevich et al. [2] 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, which is proven in an analogous manner to Theorem 3, 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.

Acknowledgements The authors would like to thank Susanne Albers, Elliot Anshelevich, Jian Li, Kirk Pruhs and Alexander Souza for interesting discussions. The work of the last author was supported by the German Research Foundation (DFG).

References

1. Albers S (2008) On the value of coordination in network design. In: SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp 294–303
2. Anshelevich E, Dasgupta A, Kleinberg JM, Tardos É, Wexler T, Roughgarden T (2004) The price of stability for network design with fair cost allocation. In: 45th Symposium on Foundations of Computer Science (FOCS), pp 295–304
3. Awerbuch B, Azar Y, Epstein A (2005) Large the price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), pp 57–66
4. Caragiannis I, Flammini M, Kaklamanis C, Kanellopoulos P, Moscardelli L (2006) Tight bounds for selfish and greedy load balancing. In: Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP), pp 311–322
5. Chen HL, Roughgarden T (2006) Network design with weighted players. In: SPAA '06: Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures, ACM, New York, NY, USA, pp 29–38, DOI <http://doi.acm.org/10.1145/1148109.1148114>
6. Christodoulou G, Koutsoupias E (2005) On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Algorithms - ESA 2005, 13th Annual European Symposium, pp 59–70
7. Christodoulou G, Koutsoupias E (2005) The price of anarchy of finite congestion games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), pp 67–73
8. Christodoulou G, Chung C, Ligett K, Pyrga E, van Stee R (2009) On the price of stability for undirected network design. In: Proceedings of the 7th Workshop on Approximation and Online Algorithms (WAOA), to appear
9. Fiat A, Kaplan H, Levy M, Olonetsky S, Shabo R (2006) On the price of stability for designing undirected networks with fair cost allocations. In: ICALP (1), pp 608–618
10. Koutsoupias E, Papadimitriou CH (1999) Worst-case equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), pp 404–413
11. Li J (2009) An $O(\log\log\log n)$ upper bound on the price of stability for undirected Shapley network design games. *Information Processing Letters*
12. Monderer D, Shapley LS (1996) Potential games. *Games and Economic Behavior* 14:124–143
13. Olonetsky S (2006) On the price of stability for designing undirected networks with fair cost allocations. Master's thesis, Tel-Aviv University
14. Papadimitriou CH (2001) Algorithms, games, and the internet. In: Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC), pp 749–753
15. Rosenthal R (1973) A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2:65–67

-
16. Roughgarden T, Tardos E (2002) How bad is selfish routing? *J ACM* 49(2):236–259
 17. Tardos E, Wexler T (2007) Network Formation Games and the Potential Function Method. In: *Algorithmic Game Theory*, Cambridge University Press, chap 19