Improved bounds on equilibria solutions in the network design game

Int J Game Theory (2018) 47:1113–1135
https://doi.org/10.1007/s00182-017-0600-z

ORIGINAL PAPER

Improved bounds on equilibria solutions in the network design game

Akaki Mamageishvili[1] · **Matúš Mihalák**[2] · **Simone Montemezzani**[3]

Accepted: 12 November 2017 / Published online: 21 November 2017

Abstract In the network design game with n players, every player chooses a path in an edge-weighted graph to connect her pair of terminals, sharing costs of the edges on her path with all other players fairly. It has been shown that the price of stability of any network design game is at most H_n, the n-th harmonic number. This bound is tight for directed graphs.
For undirected graphs, it has only recently been shown that the price of stability is at most $H_n\left(1-\frac{1}{\Theta(n^4)}\right)$, while the worst-case known example has price of stability around 2.25. We improve the upper bound considerably by showing that the price of stability is at most $H_{n/2}+\varepsilon$ for any ε starting from some suitable $n \geq n(\varepsilon)$.
We also study quality measures of different solution concepts for the multicast network design game on a ring topology. We recall from the literature a lower bound of $\frac{4}{3}$ and prove a matching upper bound for the price of stability. Therefore, we answer an open question posed by Fanelli et al. (Theor Comput Sci 562:90–100, 2015). We prove an upper bound of 2 for the ratio of the costs of a potential optimizer and of an optimum, provide a construction of a lower bound, and give a computer-assisted argument that it reaches 2 for any precision. We then turn our attention to players arriving one by one and playing myopically their best response. We provide matching lower and upper bounds of 2 for the myopic sequential price of anarchy (achieved for a worst-case order of the arrival of the players). We then initiate the study of myopic sequential price of stability and for the multicast game on the ring we construct a lower bound

✉ Akaki Mamageishvili
akaki@inf.ethz.ch

1 Department of Management, Technology and Economics, ETH Zurich, Zurich, Switzerland

2 Department of Data Science and Knowledge Engineering, Maastricht University, Maastricht, The Netherlands

3 Department of Mathematics, ETH Zurich, Zurich, Switzerland

of $\frac{4}{3}$, and provide an upper bound of $\frac{26}{19}$. To the end, we conjecture and argue that the right answer is $\frac{4}{3}$.

Keywords Network design game · Nash equilibrium · Price of stability · Ring topology · Potential-optimum price of stability/anarchy

1 Introduction

Network design game was introduced by Anshelevich et al. (2004) together with the notion of price of stability (PoS), as a formal model to study and quantify the strategic behavior of non-cooperative agents in designing communication networks. Network design game with n players is given by an edge-weighted graph G (where n does not stand for the number of vertices), and by a collection of n terminal (source-target) pairs $\{s_i, t_i\}$, $i = 1, \ldots, n$. In this game, every player i connects its terminals s_i and t_i by an s_i–t_i path P_i, and pays for each edge e on the path a fair share of its cost (i.e., all players using the edge pay the same amount totalling to the cost of the edge). A Nash equilibrium of the game is an outcome $(P_1, \ldots, P_n)$ in which no player i can pay less by changing P_i to a different path P_i'.

Nash equilibria of the network design game can be quite different from an optimal outcome that could be created by a central authority. To quantify the difference in quality of equilibria and optima, one compares the total cost of a Nash equilibrium to the cost of an optimum (with respect to the total cost). Taking the worst-case approach, one arrives at the *price of anarchy*, which is the ratio of the maximum cost of any Nash equilibrium to the cost of an optimum. The price of anarchy of network design games can be as high as n (but not higher) (Anshelevich et al. 2004). Taking the slightly less pessimistic approach leads to the notion of the *price of stability*, which is the ratio of the smallest cost of any Nash equilibrium to the cost of an optimum. The motivation behind this is that often a central authority exists, but cannot force the players into actions they do not like. Instead, a central authority can suggest to the players actions that correspond to a best Nash equilibria. Then, no player wants to deviate from the action suggested to her, and the overall cost of the outcome can be lowered (when compared to the worst case Nash equilibria).

Network design games belong to the broader class of congestion games, introduced in Rosenthal (1973), for which a function (called a *potential function*) $\Phi(P_1, \ldots, P_n)$ exists, with the property that $\Phi(P_1 \ldots, P_i, \ldots, P_n) - \Phi(P_1, \ldots, P_i', \ldots, P_n)$ exactly reflects the changes of the cost of any player i switching from P_i to P_i'. This property implies that a collection of paths $(P_1, \ldots, P_n)$ minimizing Φ necessarily needs to be a Nash equilibrium. Up to an additive constant, every congestion game has a unique potential function of a concrete form, which can be used to show that the price of stability of any network design game is at most $H_n := \sum_{i=1}^{n} \frac{1}{i}$, the n-th harmonic number, and this is tight for directed graphs (i.e., there is a network design game for which the price of stability is arbitrarily close to H_n) (Anshelevich et al. 2004).

For undirected graphs, the situation is dramatically different, obtaining tight bounds on the price of stability for undirected graphs turned out to be much more difficult. The worst case known example is an involved construction of a game by Bilò et al.

(2013) achieving in the limit the price of stability of around 2.25. While the general upper bound of H_n applies also for undirected graphs, it has not been known for a long time whether it can be any lower, until the recent work of Disser et al. (2015) who showed that the price of stability of any network design game with n players is at most $H_n \cdot \left(1 - \frac{1}{\Theta(n^4)}\right)$. Improved upper bounds have been obtained for special cases. For the case where all terminals t_i are the same, Jian (2009) showed that the price of stability is at most $O\left(\frac{\log n}{\log \log n}\right)$ (note that H_n is approximately $\ln n$). If, additionally, every vertex of the graph is a source of a player, a series of papers by Fiat et al. (2006), Lee and Ligett (2013), and Bilò et al. (2014) showed that the price of stability is in this case at most $O(\log \log n)$, $O(\log \log \log n)$, and $O(1)$, respectively. Note that, in this special case optimum solution is a minimum spanning tree and different techniques were developed for this particular structure.

Fanelli et al. (2015) restrict the graphs to be rings, and prove that the price of stability is at most $3/2$.

Further special cases concern the number of players. Interestingly, tight bounds on price of stability are known only for $n = 2$ (we do not consider the case $n = 1$ as a game) (Anshelevich et al. 2004; Christodoulou et al. 2009), while for already 3 players there are no tight bounds; for the most recent results for the case $n = 3$, see Disser et al. (2015) and Bilò and Bove (2011).

All obtained upper bounds on the price of stability use the potential function in one way or another. Our result is not an exception in that aspect. Bounding the price of stability translates effectively into bounding the cost of a best Nash equilibrium. A common approach is to bound this cost by the cost of the potential function minimizer $(P_1^{\Phi}, \ldots, P_n^{\Phi}) := \arg\min_{(P_1,\ldots,P_n)} \Phi(P_1, \ldots, P_n)$, which is (as we argued above) also a Nash equilibrium. Using just the inequality $\Phi(P_1^{\Phi}, \ldots, P_n^{\Phi}) \leq \Phi(P_1^{O}, \ldots, P_n^{O})$, where $(P_1^{O}, \ldots, P_n^{O})$ is an optimal outcome (minimizing the total cost of having all pairs of terminals connected), one obtains the original upper bound H_n on the price of stability (Anshelevich et al. 2004). In Disser et al. (2015) and Christodoulou et al. (2009) authors consider other inequalities obtained from the property that potential optimizer is also a Nash equilibrium to obtain improved upper bounds. In this paper, we consider n different specifically chosen strategy profiles $(P_1^{i}, \ldots, P_n^{i})$, $i = 1, \ldots, n$, in which players use only edges of the optimum $(P_1^{O}, \ldots, P_n^{O})$ and of the Nash equilibrium $(P_1^{\Phi}, \ldots, P_n^{\Phi})$. This idea is a generalization of the approach used by Bilò and Bove (2011) to prove an upper bound of $286/175 \approx 1.634$ for Shapley network design games with 3 players. Clearly, the potential of each of the considered strategy profile is at least the potential of $(P_1^{\Phi}, \ldots, P_n^{\Phi})$. Summing all these n inequalities and combining it with the original inequality $\Phi(P_1^{\Phi}, \ldots, P_n^{\Phi}) \leq \Phi(P_1^{O}, \ldots, P_n^{O})$ gives an asymptotic upper bound of $H_{n/2} + \varepsilon$ on the price of stability. Our result thus shows that the price of stability is strictly lower than H_n by an additive constant (namely, by $\log 2$).

Albeit the idea is simple, the analysis is not. The main difficulty is to overcome situations where socially optimum solution has an arbitrary structure. This has not been done before up to our knowledge. The analysis involves carefully chosen strategy profiles for various possible topologies of the optimum solution. These considerations can be of independent interest in further attempts to improve the bounds on the price of stability of network design games.

Additionally, equilibria minimizing potential function are regarded as stable (against noise) by Asadpour and Saberi (2009) and Alós-Ferrer and Netzer (2010), and accordingly, some authors studied the price of stability restricted to these kind of equilibria (Kawase and Makino 2013), the so-called *potential-optimum price of stability*.

As already mentioned above, one of the main motivations to study best Nash equilibria is that they can be regarded as outcomes of the game if a little coordination is present—an authority that suggests the players the strategies P_i. Then, players have no incentive to unilaterally deviate from the suggested strategy profile. It is questionable whether such an authority exists—it would need to be very strong, both computationally and imperatively. To address this applicability issue of equilibrium concepts, sequential versions of the game were studied: the players arrive one by one, and upon arrival, player i chooses *myopically* the best path P_i as if this was the end of the game (i.e., no further players would arrive). Chekuri et al. (2006) show that the total cost achieved by a worst-case permutation of the arriving players is at most $O(\sqrt{n}\log n)$ times the optimum cost. Subsequently, Charikar et al. (2008) improved this bound to $O(\log^2 n)$ [the original version (Charikar et al. 2008) is erroneous, but the authors provide corrected arguments upon request]. The worst-case approach to the order in which the players arrive naturally models the complete lack of coordination. In this paper, we suggest to study also the best-case order in which players arrive. This is motivated by the presence of an authority that can control the access to the resources over time (and thus decide an order of the arriving players). Such an authority is arguably weaker than the one mentioned above, as it does not impose any decision upon the players, and it leaves them to decide their strategies freely upon arriving. Bilò et al. (2010) studied a version of a cost sharing scheme for multicast network design game, in which each player only knows strategies of some other players, and pays fair share of edge costs that she uses based only on her information. Sequential versions described above can be modeled with this cost sharing scheme.

In the second part of the paper, we focus on one specific network topology: the ring. This is a fundamental topology in networking and communications. It is the edge-minimal topology that is resistant against a single link fault. From the decentralized point of view, *call control* comes close in spirit to network design games, in that the connecting s_i–t_i paths needs to be chosen to obey given capacities on the links (Adamy et al. 2007). The study of approximation algorithms is the counterpart to bounding the prices of anarchy and stability. Rings have also been intensively studied in the distributed setting, e.g., among plenty of others, in the context of the fundamental leader election problem (Attiya et al. 1988). In the second part of the paper we restrict ourselves to the *multicast* version in which all players share the same target vertex $t = t_i, i = 0, \ldots, n-1$ and answer the open question asked by Fanelli et al. (2015) about tight bounds of the price of stability for multicast game on a ring. We study various solutions concepts and analyze their quality compared to an optimum network (with respect to the social cost). In most cases, we are able to provide tight bounds. Furthermore, we also study the myopic sequential price of stability in general multicast network design games, and give a simpler proof of an upper bound of 4 for this class of games compared to a more general proof in Bilò et al. (2010) (cf. this with the upper bound of $\log^2 n$ on the myopic sequential price of anarchy for multicast games).

2 Preliminaries

Shapley network design game is a strategic game of n players played on an edge-weighted graph $G = (V, E)$ with non-negative edge costs c_e, $e \in E$. Each player i, $i = 1, \ldots, n$, has a *source* node s_i and a *target* node t_i. All s_i–t_i paths form the set $\mathcal{P}_i$ of the *strategies* of player i. A vector $P = (P_1, \ldots, P_n) \in \mathcal{P}_1 \times \cdots \times \mathcal{P}_n$ is called a *strategy profile*. Let $E(P) := \bigcup_{i=1}^n P_i$ be the set of all edges used in P. The *cost of player i* in a strategy profile P is $\text{cost}_i(P) = \sum_{e \in P_i} c_e / k_e(P)$, where $k_e(P) = \left|\{j | e \in P_j\}\right|$ is the number of players using edge e in P. A strategy profile $N = (N_1, \ldots, N_n)$ is a *Nash equilibrium* if no player i can unilaterally switch from her strategy N_i to a different strategy $N_i' \in \mathcal{P}_i$ and decrease her cost, i.e., $\text{cost}_i(N) \le cost_i(N_1, \ldots, N_i', \ldots, N_n)$ for every $N_i' \in \mathcal{P}_i$.

Shapley network design games are exact potential games. That is, there is a so called *potential function* $\Phi : \mathcal{P}_1 \times \cdots \times \mathcal{P}_n \to \mathbb{R}$ such that, for every strategy profile P, every player i, and every alternative strategy P_i', $cost_i(P) - cost_i(P_1, \ldots, P_i', \ldots, P_n) = \Phi(P) - \Phi(P_1, \ldots, P_i', \ldots, P_n)$. Up to an additive constant, the potential function is unique (Monderer and Shapley 1996), and is defined as

$$\Phi(P) = \sum_{e \in E(P)} \sum_{i=1}^{k_e(P)} c_e / i = \sum_{e \in E(P)} H_{k_e(P)}\, c_e.$$

To simplify the notation (e.g., to avoid writing $H_{\lceil n/2 \rceil}$), we extend H_k also for non-integer values of k by setting $H(k) := \int_0^1 \frac{1-x^k}{1-x} dx$, which is an increasing function, and which agrees with the (original) k-th harmonic number whenever k is an integer.

The *social cost* of a strategy profile P is defined as the sum of the player costs: $\text{cost}(P) = \sum_{i=1}^n \text{cost}_i(P) = \sum_{i=1}^n \sum_{e \in P_i} c_e / k_e(P) = \sum_{e \in E(P)} k_e(P)\, c_e / k_e(P) = \sum_{e \in E(P)} c_e$. A strategy profile $O(G)$ that minimizes the social cost of a game G is called a *social optimum*. Observe that the edge set of a social optimum $O(G)$ induces a forest (if there is a cycle, we could remove one of its edges without increasing the social cost). Let $\mathcal{N}(G)$ be the set of Nash equilibria of a game G. The *price of stability of a game G* is the ratio $\text{PoS}(G) = \min_{N \in \mathcal{N}(G)} \text{cost}(N)/\text{cost}(O(G))$.

Let $\mathcal{M}(G)$ be the set of Nash equilibria that are also global minimizers of the potential function Φ of the game. The *potential-optimal price of anarchy* of a game G, introduced by Kawase and Makino (2013), is defined as $\text{POPoA}(G) = \max_{N \in \mathcal{M}(G)} \text{cost}(N)/cost(O(G))$. Properties of potential optimizers were earlier observed and exploited by Asadpour and Saberi (2009) for other games.

Since $\mathcal{M}(G) \subset \mathcal{N}(G)$, it follows that $\text{PoS}(G) \le \text{POPoA}(G)$. Let $\mathcal{G}(n)$ be the set of all Shapley network design games with n players. The *price of stability of Shapley network design games* is defined as $\text{PoS}(n) = \sup_{G \in \mathcal{G}(n)} \text{PoS}(G)$. The quantity $\text{POPoA}(n)$ is defined analogously, and we get that $\text{PoS}(n) \le \text{POPoA}(n)$.
In the multicast game all t_i's are the same and we denote it by t.

Fig. 1 Multicast game on rings

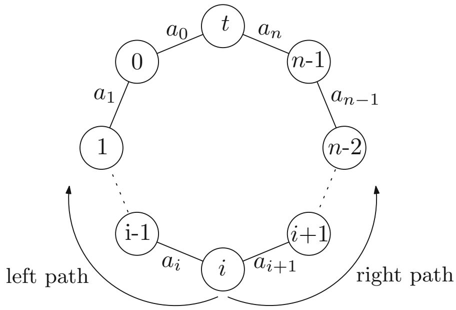

Observe also that in a multicast game an optimum network forms a Steiner tree on the terminals s_i and t. If an underlying graph G is a ring, then there are only 2 possible strategies for each player.

In the second part of the paper, we focus on the multicast game on rings. We can assume, without loss of generality, that every node but the target t is a source of exactly one player. Otherwise, we can modify the topology by the following two operations. If there are $l > 1$ players sharing the same node x of the ring as a source vertex, we make l copies of this vertex, add $l - 1$ consecutive edges of cost 0 between them to make a path of length $l - 1$, replace x in the ring with this path in a natural way, and associate each vertex with a unique source (copy of x). If there is a node x in the ring which is not a target nor a source of any player, we delete x from the ring, and connect its two neighbors by an edge of cost $c_e + c_{e'}$, where e, e' are the two adjacent edges of x. A repetitive application of these two operations preserves the cost of the optimum and Nash equilibrium strategy profiles, and also preserves the equilibrium properties of strategy profiles (if the strategies are expressed in the form "go clockwise/counterclockwise to s_i").

We label the sources (players) and the edges connecting them in a counter-clockwise order as in Fig. 1, where a_i denotes the cost of the i-th edge. Player i has exactly 2 strategies, one is to go *left*, i.e., *clockwise*, taking edges $i, i-1, \ldots 0$, or to go *right*, i.e., *counterclockwise*, taking edges $i+1, \ldots n$. Note that here players (source nodes) are 0-indexed unlike in the general setting. Observe that the optimum strategy profile is the one which uses all edges except the most expensive edge. Let o denote the most expensive edge. Then the (social) cost of an optimum network is $\sum_{i \neq o} a_i$.

The myopic sequential price of anarchy/stability is the worst-case/best-case ratio of the costs of a strategy profile that can be obtained by ordering the players as in a permutation π and letting player $\pi(i)$ choose the best-response $p_{\pi(i)}$ in the game induced by the first i players $\pi(1), \pi(2), \ldots, \pi(i)$ and of an optimum profile.

Note on related concepts. The term *sequential price of anarchy* has been used (Leme et al. 2012; Angelucci et al. 2015) to express a different, yet still closely related, concept compared to the notion of the myopic sequential price of anarchy/stability. In the sequential price of anarchy, players also come one by one, and decide their strategy upon arrival, but the stability of the outcome is measured in terms of subgame perfect equilibria. In some sense, the game resembles extensive games. Observe that profiles

p that get compared to optima in the myopic sequential price of anarchy/stability are in general no Nash equilibria.

3 The $\approx H_{n/2}$ upper bound

The main result of the paper is the new upper bound on the price of stability, as stated in the following theorem.

Theorem 3.1 *$PoS(n) \leq H_{n/2} + \varepsilon$, for any $\varepsilon > 0$ given that $n \geq n(\varepsilon)$ for some suitable $n(\varepsilon)$.*

We consider a Nash equilibrium N that minimizes the potential function Φ. For each player i we construct a strategy profile S^i as follows. Every player $j \neq i$, whenever possible (the terminals of players i and j lie in the same connected component of the optimum O), uses edges of $E(O(G))$ to reach s_i, from there it uses the Nash equilibrium strategy (a path) of player i to reach t_i, and from there it again uses edges of $E(O(G))$ to reach the player j's other terminal node. From the definition of N, we then obtain the inequality $\Phi(N) \leq \Phi(S^i)$. We then combine these n inequalities in a particular way with the inequality $\Phi(N) \leq \Phi(O(G))$, and obtain the claimed upper bound on the cost of N.

The proof of Theorem 3.1 is structured in the following way. We first prove the theorem for the special case where an optimum $O(G)$ contains an edge that is used by every player. We then extend the proof of this special case, first to the case where $E(O(G))$ is a tree, but with no edge used by every player, and, second, to the case where $E(O(G))$ is a general forest (i.e., not one connected component).

We will use the following notation. For a strategy profile $P = (P_1, \dots, P_n)$ and a set $U \subset \{1, \dots, n\}$, we denote by P_U the set of edges $e \in E$ for which $\{j | e \in P_j\} = U$ and by P^l the set of edges $e \in E$ for which $|\{j | e \in P_j\}| = l$. That is, P_U is the set of edges used in P by exactly the players U, and $P^l = \bigcup_{\substack{U \subset \{1,\dots,n\} \\ |U|=l}} P_U$ is the set of edges used by exactly l many players. Then the edges used by player i in P are $\bigcup_{\substack{U \subset \{1,\dots,n\} \\ i \in U}} P_U$. We stress that for every player $i \in U$, the edges of P_U are part of the strategy P_i; this implies that, whenever $E(P)$ induces a forest, the source s_i and the target t_i are in two different connected components of $E(P) \setminus P_U$. For any set of edges $F \subset E$, let $|F| := \sum_{e \in F} c_e$. We then have, for instance, that the cost of player i in P is given by $\text{cost}_i(P) = \sum_{\substack{U \subset \{1,\dots,n\} \\ i \in U}} \frac{|P_U|}{|U|}$.

From now on, G is an arbitrary Shapley network design game with n players, $N = (N_1, \dots, N_n)$ is a Nash equilibrium minimizing the potential function and $O = (O_1, \dots, O_n)$ is an arbitrary social optimum so that $E(O)$ has no cycles.

3.1 Case O^n is not empty

In this section we assume that O^n is not empty. In this case, $E(O)$ is actually a tree. Then, $E(O) \setminus O^n$ is formed by two disconnected trees, which we call O^- and O^+, such that each player has the source node in one tree and the target node in the other

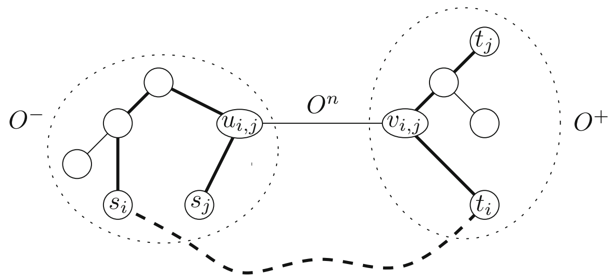

Fig. 2 The non dashed lines are the edges of $E(O)$, the dashed line is the Nash strategy N_i. The path S^i_j from s_j to t_j is given by the thicker dashed and non dashed lines

tree (see also Fig. 2). Without loss of generality, assume that all source nodes s_i are in O^-. Given two players i and j, let $u_{i,j}$ be the first[1] edge of $O_i \cap O_j$ and $v_{i,j}$ be the last edge of $O_i \cap O_j$. Notice that every edge between s_i and $u_{i,j}$ is used in O by player i but not by player j. That is, each edge e between s_i and $u_{i,j}$ satisfies $e \in O_i$ and $e \notin O_j$, or equivalently, $e \in \bigcup_{\substack{U\subset\{1,\ldots,n\}\\ i\in U, j\notin U}} O_U$. An analogous statement holds for each edge e between t_i and $v_{i,j}$.

For every player i, we define a strategy profile S^i, where player $j = 1, \ldots, n$ uses the following s_j–t_j path S^i_j (see Fig. 2 for an example):

1. From s_j to $u_{i,j}$, it uses edges of O^-.
2. From $u_{i,j}$ to s_i, it uses edges of O^-.
3. From s_i to t_i, it uses edges of N_i.
4. From t_i to $v_{i,j}$, it uses edges of O^+.
5. From $v_{i,j}$ to t_j, it uses edges of O^+.

If S^i_j contains cycles, we skip them to obtain a simple path from s_j to t_j. This can be the case if N_i is not disjoint from $E(O)$, so that an edge appears both in step 3 and in one of the steps 1, 2, 4 or 5. Observe that the path S^i_j uses exactly the edges of O_U for $i \in U$, $j \notin U$ (in steps 2 and 4), the edges of O_U for $i \notin U$, $j \in U$ (in steps 1 and 5) and the edges of N_U for $i \in U$ (in step 3). We now can prove the following lemma.

Lemma 3.2 *For every* $i \in \{1, \ldots, n\}$,

$$\Phi(N) \le \Phi(S^i) \le \sum_{\substack{U\subset\{1,\ldots,n\}\\ i\in U}} H_n|N_U| + \sum_{\substack{U\subset\{1,\ldots,n\}\\ i\in U}} H_{n-|U|}|O_U| + \sum_{\substack{U\subset\{1,\ldots,n\}\\ i\notin U}} H_{|U|}|O_U|. \tag{1}$$

Proof The first inequality of (1) holds because, by assumption, N is a global minimum of the potential function Φ.

To prove the second inequality, recall that for any strategy profile P we can write $\Phi(P) = \sum_{e\in P} H_{k_e(P)}c_e = \sum_{U\subset\{1,\ldots,n\}} H_{|U|}|P_U|$. In our case, every edge $e \in S^i$ belongs either to N_U, $U \subset \{1, \ldots, n\}$, $i \in U$, or to O_U, and we therefore sum only

[1] The edges are ordered naturally along the path from s_i to t_i.

over these terms. We now show that, in our sum, the cost c_e of every edge e in S^i is accounted for with at least coefficient $H_{k_e(S^i)}$.

For the first sum in the right hand side of (1), obviously at most n players can use an edge of N_U, $i \in U$, i.e., $k_e(S^i) \leq n$. To explain the second and third sums, notice that if an edge $e \in O_U$ that is present in S^i also belongs to N_i, its cost is already accounted for in the first sum. So, we just have to look at edges that are only present in steps 1, 2, 4 and 5 of the definition of S^i_j.

To explain the second sum, let $i \in U$. Then, as we already noted, in the definition of S^i_j, player j uses edges of O_U with $i \in U$ only if $j \notin U$ (in steps 2 and 4). Since there are exactly $n - |U|$ players that satisfy $j \notin U$, this explains the second sum.

Finally, to explain the third sum, let $i \notin U$. Similarly to the previous argument, in the definition of S^i_j, player j uses edges of O_U with $i \notin U$ only if $j \in U$ (in steps 1 and 5). Since there are exactly $|U|$ players that satisfy $j \in U$, this explains the third sum. □

We now show how to combine Lemma 3.2 with the inequality $\Phi(N) \leq \Phi(O)$ to prove Theorem 3.1, whenever $O^n \neq \emptyset$.

Lemma 3.3 *Suppose that inequality* (1) *holds for every* i. *Then, for* $x = \frac{n-H_n}{H_n-1}$,

$$PoS(G) \leq \frac{n+x}{n+x-H_n} H_{\frac{n+x}{2}} \leq H_{n/2} + \varepsilon$$

holds for any $\varepsilon > 0$, *given that* $n \geq n(\varepsilon)$ *for some suitable* $n(\varepsilon)$.

Proof We sum (1) for $i = 1, \ldots, n$ to obtain

$$\begin{aligned} n\Phi(N) &\leq \sum_{i=1}^{n} \left(\sum_{\substack{U\subset\{1,\ldots,n\}\\ i\in U}} H_n|N_U| + \sum_{\substack{U\subset\{1,\ldots,n\}\\ i\in U}} H_{n-|U|}|O_U| + \sum_{\substack{U\subset\{1,\ldots,n\}\\ i\notin U}} H_{|U|}|O_U| \right) \\ &= \sum_{U\subset\{1,\ldots,n\}} |U|H_n|N_U| + \sum_{U\subset\{1,\ldots,n\}} |U|H_{n-|U|}|O_U| \\ &\quad + \sum_{U\subset\{1,\ldots,n\}} (n-|U|)H_{|U|}|O_U| = \sum_{l=1}^{n} lH_n|N^l| \\ &\quad + \sum_{l=1}^{n} (lH_{n-l} + (n-l)H_l)|O^l|. \end{aligned}$$

Since $\Phi(N) = \sum_{l=1}^{n} H_l|N^l|$, by putting all terms relating to N on the left hand side we obtain

$$\sum_{l=1}^{n} (nH_l - lH_n)|N^l| \leq \sum_{l=1}^{n} (lH_{n-l} + (n-l)H_l)|O^l|. \tag{2}$$

On the other hand, we have $\Phi(N) \leq \Phi(O)$, which we can write as

$$\sum_{l=1}^{n} H_l |N^l| \leq \sum_{l=1}^{n} H_l |O^l|. \tag{3}$$

If we multiply (3) by $x = \frac{n-H_n}{H_n-1}$ and sum it with (2) we get

$$\sum_{l=1}^{n} ((n+x)H_l - lH_n)|N^l| \leq \sum_{l=1}^{n} (lH_{n-l} + ((n+x)-l)H_l)|O^l|. \tag{4}$$

Let $\alpha(l) = (n+x)H_l - lH_n$ and $\beta(l) = lH_{n-l} + ((n+x)-l)H_l$. We will show that $\min_{l\in\{1,\ldots,n\}} \alpha(l) = n + x - H_n$ and that $\max_{l\in\{1,\ldots,n\}} \beta(l) \leq (n+x)H_{\frac{n+x}{2}}$. This will allow us to bound the left and right hand side of (4), giving us the desired bound on the price of stability.

To prove $\min_{l\in\{1,\ldots,n\}} \alpha(l) = n + x - H_n$, we observe that $\alpha(l)$ first increases and then decreases and that $\alpha(1) = \alpha(n)$. By the choice of x the values at the two extremes coincide, the minimum is $\alpha(1) = n + x - H_n$ by inserting 1 in the formula of $\alpha(l)$.

To prove $\max_{l\in\{1,\ldots,n\}} \beta(l) \leq (n+x)H_{\frac{n+x}{2}}$, we first show that $\theta(l) = lH_{n-l} + (n-l)H_l$ has maximum $nH_{n/2}$. Since θ is symmetric around $n/2$, we just have to show that the difference $\theta(l+1) - \theta(l)$ is always positive for $l+1 \leq n/2$. This proves that θ reaches at $l = n/2$ the maximum value of $\frac{n}{2}H_{n/2} + \frac{n}{2}H_{n/2} = nH_{n/2}$. We have that

$$\begin{aligned}
\theta(l+1) - \theta(l) &= (l+1)H_{n-(l+1)} + (n-(l+1))H_{l+1} - (lH_{n-l} + (n-l)H_l) \\
&= lH_{n-l} + H_{n-l} - \frac{l+1}{n-(l+1)} + (n-l)H_l - H_l + \frac{n-(l+1)}{l+1} \\
&\quad - lH_{n-l} - (n-l)H_l = \frac{n-(l+1)}{l+1} - \frac{l+1}{n-(l+1)} + H_{n-l} - H_l.
\end{aligned}$$

The term $\frac{n-(l+1)}{l+1} - \frac{l+1}{n-(l+1)}$ is positive if $n-(l+1) \geq l+1$, that is if $l+1 \leq n/2$. Since H is an increasing function, $H_{n-l} - H_l$ is positive if $l \leq n/2$, in particular if $l + 1 \leq n/2$. This proves our claim that $\theta(l) = lH_{n-l} + (n-l)H_l$ has maximum $nH_{n/2}$.

Since H is an increasing function, we then have the bound

$$\beta(l) = lH_{n-l} + ((n+x)-l)H_l \leq lH_{(n+x)-l} + ((n+x)-l)H_l \leq (n+x)H_{\frac{n+x}{2}}.$$

We can now finally prove Lemma 3.3. We know that

$$(n+x-H_n)\,\mathrm{cost}(N) = (n+x-H_n)\sum_{l=1}^{n}|N^l| \leq \sum_{l=1}^{n}((n+x)H_l - lH_n)|N^l|,$$

$$\sum_{l=1}^{n}(lH_{n-l} + ((n+x)-l)H_l)|O^l| \leq (n+x)H_{\frac{n+x}{2}} \sum_{l=1}^{n}|O^l| = (n+x)H_{\frac{n+x}{2}}\,\mathrm{cost}(O),$$

which together with (4) proves that $\mathrm{PoS}(G) \le \frac{\mathrm{cost}(N)}{\mathrm{cost}(O)} \le \frac{n+x}{n+x-H_n} H_{\frac{n+x}{2}}$.

Now observe that for any ε there is an $n(\varepsilon)$ so that $\frac{n+x}{n+x-H_n} H_{\frac{n+x}{2}} \le H_{n/2} + \varepsilon$ whenever $n \ge n(\varepsilon)$ because of the following: $\frac{n+x}{n+x-H_n} H_{\frac{n+x}{2}} = H_{\frac{n+x}{2}} + \frac{H_n}{n+x-H_n} H_{\frac{n+x}{2}} = H_{\frac{n+x}{2}} + O(H_n^2/n) \le H_{\frac{n+x}{2}} + \frac{\varepsilon}{2}$. Then again for large enough n we get $H_{\frac{n+x}{2}} \le H_{\frac{n}{2}} + (\frac{1}{n/2} \cdot (x/2)) = H_{n/2} + O(1/H_n) \le H_{n/2} + \varepsilon/2$. □

3.2 Case O^n is empty

In the previous section we proved Theorem 3.1 if $O^n \neq \emptyset$ by constructing for every pair of players i and j a particular path S_j^i that uses edges of $E(O)$ to go from s_j to s_i and from t_j to t_i.

If $E(O)$ is not connected, then there is a pair of players i, j for which s_i and s_j are in different connected components of $E(O)$, and we cannot define the path S_j^i. Even if $E(O)$ is connected, but $O^n = \emptyset$, there might be a pair of players i and j for which the path S_j^i exists, but this path is not optimal. See Fig. 7 for an example: the path S_j^i (before cycles are removed to make S_j^i a simple path) traverses some edges of $E(O)$ twice, including the edge denoted by e in the figure. The same holds even if we exchange the labeling of s_i and t_i. Thus, we may need to define a new path T_j^i for some players i and j.

To define the new path T_j^i, let us introduce some notation. Given two players i, j and two nodes $x_i \in \{s_i, t_i\}$, $x_j \in \{s_j, t_j\}$ in the same connected component of $E(O)$, let $O(x_i, x_j)$ be the unique path in $E(O)$ between x_i and x_j. If s_i and s_j are in the same connected component of $E(O)$, let $(T_j^i)'$ (respectively $(T_j^i)''$) be the following s_j–t_j path:

1′. From s_j to s_i (respectively t_i), it uses edges of $O(s_i, s_j)$ (respectively $O(t_i, s_j)$).
2′. From s_i (respectively t_i) to t_i (respectively s_i), it uses edges of N_i.
3′. From t_i (respectively s_i) to t_j, it uses edges of $O(t_i, t_j)$ (respectively $O(s_i, t_j)$).

If $(T_j^i)'$ or $(T_j^i)''$ contain cycles, we skip them to obtain a simple path from s_j to t_j. See Fig. 3 for an example of $(T_j^i)'$ and Fig. 5 for an example of $(T_j^i)''$.

Notice that in the previous section, we had $S_j^i = (T_j^i)'$ (where steps 1 and 2 are now step 1′; steps 4 and 5 are now step 3′) and $O(s_i, s_j) \cap O(t_i, t_j) = \emptyset$, since $O(s_i, s_j) \subset O^-$ and $O(t_i, t_j) \subset O^+$. This ensured that there was no edge that is traversed both in step 1′ and 3′, which would make Lemma 3.2 not hold. In general, $O(s_i, s_j) \cap O(t_i, t_j) = \emptyset$ does not have to hold; for example in Fig. 7 we have $e \in O(s_i, s_j) \cap O(t_i, t_j)$. We call the path $(T_j^i)'$ (respectively $(T_j^i)''$) *O-cycle free* if $O(s_i, s_j) \cap O(t_i, t_j) = \emptyset$ (respectively if $O(s_i, t_j) \cap O(t_i, s_j) = \emptyset$). For instance, in Fig. 7 both $(T_j^i)'$ and $(T_j^i)''$ are not *O-cycle free*.

We are now ready to define the path T_j^i for two players i and j. If s_i and s_j are in the same connected component of $E(O)$, we set $T_j^i = (T_j^i)'$ (respectively $T_j^i = (T_j^i)''$) if $(T_j^i)'$ (respectively $(T_j^i)''$) is *O-cycle free*. Otherwise, we set $T_j^i = O_j$. Similar to the previous section, let $T^i = (T_1^i, \ldots, T_n^i)$. That is, in T^i a player j uses the optimal path

O_j if the paths $(T_j^i)'$ and $(T_j^i)''$ are not defined (meaning that s_i and s_j are in different connected components of $E(O)$), or if they are not O*-cycle free* (meaning that they use some edges of $E(O)$ twice). Otherwise, player j uses the O*-cycle free* path.

The following lemma shows that the paths T^i satisfy the requirements of Lemma 3.3 if $E(O)$ is connected but $O^n = \emptyset$. A subsequent lemma will then show that the requirements of Lemma 3.3 are satisfied even if $E(O)$ is not connected.

Lemma 3.4 *If $E(O)$ is connected, then for every $i \in \{1, \ldots, n\}$*

$$\Phi(N) \leq \Phi(T^i) \leq \sum_{\substack{U \subset \{1,\ldots,n\} \\ i \in U}} H_n |N_U| + \sum_{\substack{U \subset \{1,\ldots,n\} \\ i \in U}} H_{o_i(U)} |O_U| + \sum_{\substack{U \subset \{1,\ldots,n\} \\ i \notin U}} H_{|U|} |O_U|, \tag{5}$$

with $o_i(U) \leq n - |U|$.

Proof Since the initial part of the proof is exactly the same as the proof of Lemma 3.2, we only prove that the cost c_e of every edge e in T^i is accounted for with at least coefficient $H_{k_e(T^i)}$ in the right hand side of (5). In particular, we just look at edges that are only present in steps $1'$ and $3'$ of the definition of T_j^i, since an edge $e \in O_U$ that also belongs to N_i has its cost already accounted for in the first sum.

To explain the second and third sum, let $U \subset \{1, \ldots, n\}$ and $e \in O_U$. We will look at all the possibilities of where the nodes s_i, s_j, t_i and t_j can be in the tree $E(O)$ and see whether e can be traversed in the path T_j^i. Denote by e^- and e^+ the two distinct connected components of $E(O) \backslash \{e\}$. Then, by the definition of O_U, each player $k \in U$ has $s_k \in e^-$ and $t_k \in e^+$, or viceversa. Always by the definition of O_U, each player $k \notin U$ has either $s_k, t_k \in e^-$ or $s_k, t_k \in e^+$.

To explain the third sum of (5), let $i \notin U$. For illustration purposes, assume without loss of generality that $s_i, t_i \in e^-$. Then, the only possibilities are that

- $j \in U$. Then e can be traversed, since T_j^i has to go from e^- to e^+ to connect s_j and t_j. See Fig. 3 for an illustration in the case $T_j^i \neq O_j$ and Fig. 4 for the case $T_j^i = O_j$.
- $j \notin U, s_j, t_j \in e^-$. Then e cannot be traversed, since all terminal nodes are in e^- and there is no need to traverse e. See Fig. 5 for an illustration in the case $T_j^i \neq O_j$ and Fig. 6 for the case $T_j^i = O_j$.
- $j \notin U, s_j, t_j \in e^+$. Then e cannot be traversed, since both $(T_j^i)'$ and $(T_j^i)''$ traverse e twice, so we must have $T_j^i = O_j$. See Fig. 7 for an illustration.

As we can see, e can be traversed only if $j \in U$, that is, at most $|U|$ times. This explains the third sum of (5).

Finally, to explain the second sum of (5), let $i \in U$. The only possibilities are that

- $j \in U$. Then e cannot be traversed, since at least one of $(T_j^i)'$ or $(T_j^i)''$ is a O*-cycle free* path that does not traverse e. See Fig. 8 for an illustration.
- $j \notin U$ and $T_j^i \neq O_j$. Then e can be traversed, since s_j and t_j are in the same connected component of $E(O) \backslash \{e\}$, but s_i and t_i are in different ones. See Fig. 9 for an illustration.

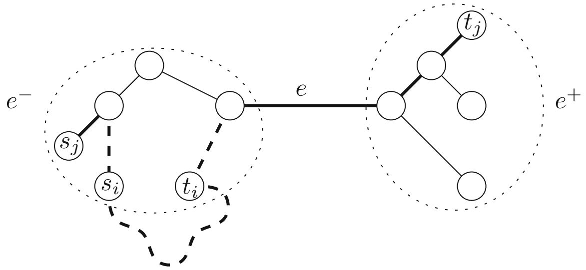

Fig. 3 $i \notin U$, $j \in U$ and $T_j^i \neq O_j$. Then e can be traversed in the path T_j^i

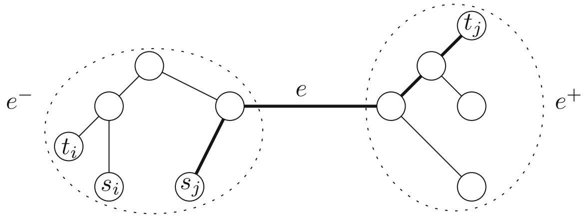

Fig. 4 $i \notin U$, $j \in U$ and $T_j^i = O_j$. Then e can be traversed in the path T_j^i

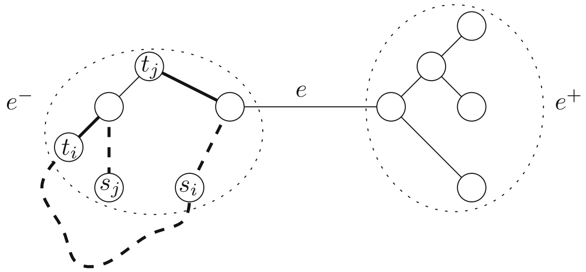

Fig. 5 $i \notin U$, $j \notin U$, $s_j, t_j \in e^-$ and $T_j^i \neq O_j$. Then e cannot be traversed in the path T_j^i

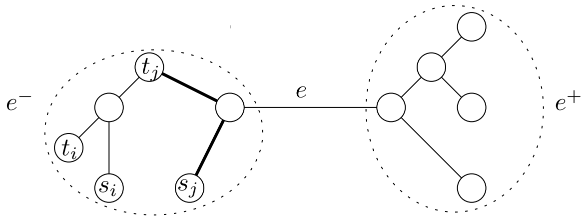

Fig. 6 $i \notin U$, $j \notin U$, $s_j, t_j \in e^-$ and $T_j^i = O_j$. Then e cannot be traversed in the path T_j^i

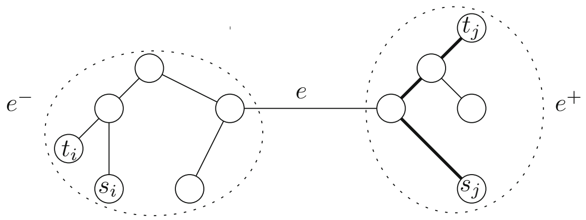

Fig. 7 $i \notin U$, $j \notin U$, $s_j, t_j \in e^+$ and $T_j^i = O_j$. Then e cannot be traversed in the path T_j^i

Fig. 8 $i \in U$, $j \in U$ and $T_j^i \neq O_j$. Then e cannot be traversed in the path T_j^i

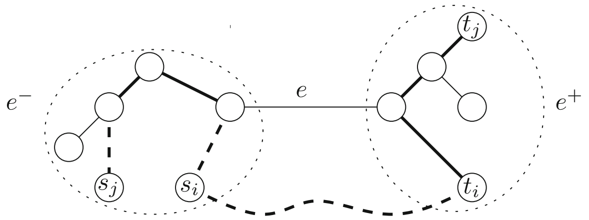

Fig. 9 $i \in U$, $j \notin U$, and $T_j^i \neq O_j$. Then e can be traversed in the path T_j^i

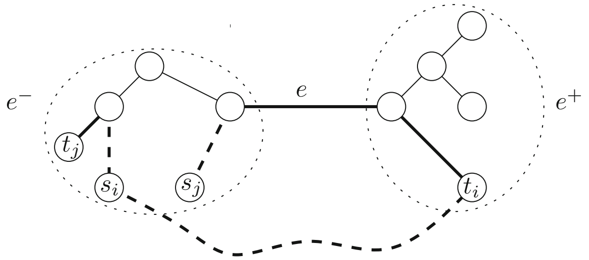

Fig. 10 $i \in U$, $j \notin U$, and $T_j^i = O_j$. Then e cannot be traversed in the path T_j^i

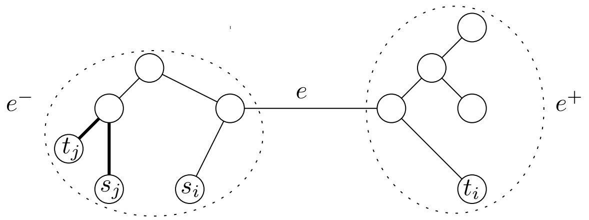

- $j \notin U$ and $T_j^i = O_j$. Then e cannot be traversed, since s_j and t_j are in the same connected component of $E(O) \backslash \{e\}$ and we just take the direct path between them, which does not traverse e. See Fig. 10 for an illustration.

Let $o_i(U)$ be the number of $j \notin U$ with $T_j^i \neq O_j$. Then, as we can see, e is traversed at most $o_i(U) \leq n - |U|$ times. This explain the second sum of (5) and finishes the proof of Lemma 3.4. □

Theorem 3.1 follows directly if $E(O)$ is connected but O^n is empty by Lemmas 3.4 and 3.3. The following lemma handles the last case we have left to analyze, which is when $E(O)$ is not a connected tree. This, together with Lemma 3.3, finishes the proof of Theorem 3.1.

Lemma 3.5 *Let $E(O) = C_1 \sqcup \cdots \sqcup C_q$, with each C_m being a connected component of $E(O)$. Let R_m be the set of players j with $s_j, t_j \in C_m$. Then for a player $i \in R_k$*

$$\Phi(N) \leq \Phi(T^i) \leq \sum_{\substack{U \subset \{1,\ldots,n\} \\ i \in U}} H_n |N_U| + \sum_{\substack{U \subset R_k \\ i \in U}} H_{o_i(U)} |O_U| + \sum_{\substack{U \subset R_m \text{ for some } m \\ i \notin U}} H_{|U|} |O_U|, \tag{6}$$

with $o_i(U) \leq |T_k| - |U| \leq n - |U|$.

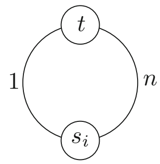

Fig. 11 In the worst equilibrium, all players use the edge of cost n

Proof Since the initial part of the proof is exactly the same as the proof of Lemmas 3.2 and 3.4, we only prove that the cost c_e of every edge e in T^i is accounted for with at least coefficient $H_{k_e(T^i)}$ in the right hand side of (6). In particular, we just look at edges that are only present in steps $1'$ and $3'$ of the definition of T^i_j, since an edge $e \in O_U$ that also belongs to N_i has its cost already accounted for in the first sum.

To explain the second and third sum, let $U \subset \{1, \ldots, n\}$ and $e \in O_U$. Notice that if $U \not\subset R_m$ for every m, then O_U is the empty set and e does not contribute anything to $\Phi(T^i)$. We begin by looking at the second sum.

Notice that since $i \in R_k$, the only possibility to have $i \in U$ is that $U \subset R_k$. By the definition of T^i the players $j \in R_m, m \neq k$ use the path O_j, which does not traverse e. With the exact same reasoning of Lemma 3.4, by looking at all the possibilities of where s_i, t_i, s_j and t_j can be in C_k, we can see that e can be traversed by player $j \in R_k$ only if $j \notin U$ and $T^i \neq O_j$. If we then define the number of players $j \in T_k$ with this property to be $o_i(U) \leq |T_k| - |U| \leq n - |U|$, the second sum in the right hand side of (6) is explained.

Finally, for the third sum, we fix $i \notin U$ and look at the cases $U \subset R_k$ and $U \subset R_m$, $m \neq k$ separately.

Suppose first that $U \subset R_k$. By the definition of T^i the players $j \in R_m, m \neq k$ use the path O_j, which does not traverse e. With the exact same reasoning of Lemma 3.4, by looking at all the possibilities of where s_i, t_i, s_j and t_j can be in C_k, we can see that e can be traversed by player $j \in R_k$ only if $j \in U$. That is, by at most $|U|$ players. This explains the third sum for the case $U \subset R_k$.

We now look at the case $U \subset R_m, m \neq k$. By the definition of T^i, players $j \in R_l$, $l \neq m$ do not traverse e, since they only use edges of C_l (if $l \neq k$) or edges of C_k and of N_i (if $l = k$). Players $j \in R_m$ use the path O_j, and by the definition of O_U exactly $|U|$ players traverse e. This explains the third sum for the case $U \subset R_m, m \neq k$, which finishes the proof. □

4 Price of anarchy/stability for multicast on rings

It is known that the price of anarchy on general graphs is at most n, and that this bound is tight. The tight example actually is a multicast game on a ring, and the general analysis of the price of anarchy thus carries over to our multicast game on rings. For completeness, we show the example in Fig. 11.

Theorem 4.1 (Anshelevich et al. 2004) *The price of anarchy for mutlicast games on rings is at most n. This is tight.*

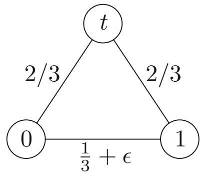

Fig. 12 Example of a lower bound 4/3

We now turn our attention to the price of stability. The example from Fig. 12, due to Anshelevich et al. (2004), shows that the price of stability can be as high as 4/3 (observe that the game possesses a unique Nash equilibrium where both players use the direct edge to get connected to t). We now show that the price of stability cannot get larger than that for multicast games on rings, and therefore answer the open question asked by Fanelli et al. (2015).

Theorem 4.2 *The price of stability in the multicast game on rings is at most* $\frac{4}{3}$.

In the proof of the theorem we will use the following lemma.

Lemma 4.3 *If a strategy profile* P *in which an edge* i *is not used is not Nash equilibrium, then either player* i *or player* $i-1$ *can improve her cost by changing her strategy.*

Proof Since the strategy profile P is not a Nash equilibrium, there exists a player k that can change her strategy and improve the cost. Assume, without loss of generality, that $k < i-1$. Since edge i is not used in P, it follows that player k uses the left path to get to t. The cost of k in P is thus $\sum_{l=0}^{k} \frac{a_l}{i-l}$, which is, by our assumption, bigger than the cost of k if she switches to the right path, i.e., bigger than $\sum_{l=k+1}^{i-1} \frac{a_l}{i-l+1} + \sum_{l=i}^{n} \frac{a_l}{l-i+1}$. It follows that player $i-1$ also uses the left path in P, and thus her cost is at least the cost of player k, whereas the alternative cost of $i-1$ if she switches to the right path is at most the alternative cost of player k. Hence, the alternative cost of player $i-1$ is smaller than her cost in P, and player $i-1$ thus improves her cost as well. □

Proof of Theorem 4.2 Consider an optimum strategy profile and let o be the edge that is not used in it. If the optimum is also Nash equilibrium, then price of stability is 1 and the claim follows. Otherwise, the optimum is not a Nash equilibrium and, by Lemma 4.3, one of the endpoints of the edge o can improve its cost. Assume, without loss of generality, that player $o-1$ can improve. We now consider the following best-response dynamics: let $o-1$ improve; then, edge $o-1$ is not used, and in case we have not reached Nash equilibrium, let player $o-2$ improve (the player $o-2$ must be able to improve by Lemma 4.3), and so on, until some player $o-k$ cannot improve anymore (this happens at the latest for player 0), and we reach a Nash equilibrium.

We will show that the social cost of a Nash equilibrium that is reached by this best response dynamics is maximized for $k=1$, i.e., for the strategy profile reached after one step of the dynamics. We then show that the cost of such a profile is at most $4/3$ times the cost of the optimum, which proves the theorem.

Let us first show the second part. Assume therefore that player $o-1$ switches to improve her cost, and the resulting profile is an equilibrium. In particular, we have that player $o-2$ does not want to switch. This can be expressed by the following two inequalities: $\sum_{l=o}^{n} \frac{a_l}{l-o+1} \leq \sum_{l=0}^{o-1} \frac{a_l}{o-l}$, and $\sum_{l=0}^{o-2} \frac{a_l}{o-1-l} \leq \sum_{l=o-1}^{n} \frac{a_l}{l-o+2}$. We further introduce a normalization of the edge costs so that the edges in the optimum sum up to 1. Thus, we obtain the *normalization* equation $\sum_{i=0,i\neq o}^{n} a_i = 1$. Now, taking the first inequality with weight 5, the second with weight 1, and the normalization equality with weight 6, we obtain that the cost of the Nash equilibrium where edge $o-1$ is not used has cost $\sum_{i=0,i\neq o-1}^{n} a_i$ at most $\frac{4}{3}$.

We can proceed in the same way for every other value of $k = 2, 3, \ldots$ for which the reached Nash equilibrium does not use edge $o-k$. For every k, we get for each of the players $o-k-1, o-k, \ldots, o-1$ an inequality stating that the player did not want, respectively wanted to swap her strategy. For all values of $k = 1, 2, 3, 4, 5, 6, 7$, we provide in the appendix the coefficients with which we need to take the inequalities and to obtain the upper bound of at most 4/3 on the cost of the Nash equilibrium.

If the length of the best-response dynamics is 8 or more, it follows that we do not need to add further inequalities, and the 7 inequalities obtained for the first 7 deviating players are enough to show the upper bound of 4/3 on the cost of the reached Nash equilibrium. □

5 Potential-optimum price of anarchy for multicast on rings

The potential-optimum price of anarchy/stability has been first studied, in the context of the network design games, by Kawase and Makino (2013). Besides other results, they proved that for multicast network design games, the two values collide. Therefore, in the following, we only study the potential-optimum price of anarchy (POPoA for short), and we show that it is at most two for rings, and provide an infinite family of examples with increasing POPoA, which we conjecture converges to two, but leave the formal analysis as an open problem. We have analyzed one such game from the family which shows that POPoA can be as large as 1.99992.

Theorem 5.1 *POPoA is at most* 2 *in the multicast game on rings.*

Proof Consider an optimal strategy profile O and let o be the edge that is not used in it. Consider a potential optimum strategy profile P and let p be the edge in it that is not used by any player. Assume, without loss of generality, that $p < o$.

By the definition of P, we have, for any strategy profile Q, $\Phi(P) \leq \Phi(Q)$, and in particular $\Phi(P) \leq \Phi(O)$, i.e.,

$$\sum_{i=0}^{p-1} a_i \cdot H_{p-i} + \sum_{i=p+1}^{n} a_i \cdot H_{i-p} \leq \sum_{i=0}^{o-1} a_i \cdot H_{o-i} + \sum_{i=o+1}^{n} a_i \cdot H_{i-o}. \quad (7)$$

We now concentrate on a_o and show that a_o is at most the cost of the optimum, i.e., at most $\sum_{i\neq o} a_i$. This then shows that any strategy profile (and, in particular, P) has cost at most twice the cost of the optimum.

Isolate in the second sum of the left hand side (LHS for short) of Eq. (7) the term with a_o and put the rest of the sum to the right hand side (RHS). This rest will dominate the second sum on the RHS, and by neglecting the resulting negative number, we get that $\sum_{i=0}^{p-1} a_i \cdot H_{p-i} + a_o \cdot H_{o-p} \leq \sum_{i=0}^{o-1} a_i \cdot H_{o-i}$, or, equivalently, that $a_o \leq \frac{\sum_{i=0}^{o-1} a_i \cdot H_{o-i} - \sum_{i=0}^{p-1} a_i \cdot H_{p-i}}{H_{o-p}}$. Consider the coefficients c_i for each a_i, $i = 0, ..., o-1$ and rewrite the right hand side of the latter inequality as $\sum_{i=0}^{o-1} c_i a_i$. Then if $i < p$, $c_i = \frac{H_{o-i} - H_{p-i}}{H_{o-p}} \leq 1$, while if $i \geq p$ then $c_i = \frac{H_{o-i}}{H_{o-p}} \leq 1$. Thus $a_o \leq \sum_{i=0}^{o-1} c_i a_i \leq \sum_{i=0}^{o-1} a_i \leq \sum_{i \neq o} a_i$, which proves the claim and thus the theorem. □

We now provide a construction of a game which shows that POPoA is at least 1.99992. We conjecture that the construction can be used to prove an asymptotic lower bound of 2 on POPoA.

Consider $2l + 1$ non-negative variables $a_0, \ldots, a_{2 \cdot l}$ that sum up to 1, where l is some constant. Suppose n is sufficiently large number, $o = n$, $p = l - 1$ and a_n is equal to $\frac{H_n - a}{H_n}$, for some constant a. Note that since n tends to infinity, a_o tends to 1 and therefore, optimal strategy profile leaves out the edge with number $o = n$. We assume that $a_i = 0$ for each $n > i > 2l$. We construct a linear program for obtaining the lower bound on the potential optimum price of anarchy. Constraints of the linear program come from the comparison of the potentials of the strategy profiles which do not use edge i for $i = 0, \ldots, i = 2l$ to the potential of P (the strategy profile minimizing Φ) that does not include the p-th edge of cost a_p. The variables of the linear program are a_i's. Note that after canceling the coefficients of a_i's on both sides in the linear program constraints, the coefficient in front of a_n is a sum of a constant number of terms converging to 0 for n tending to infinity, so these terms can be neglected. The sequence of potentials of the strategy profiles which do not use edge i for $n > i > 2l$ first increases when i increases and then decreases towards n. For this reason we only need to consider $2l + 1$ constraints. The aim is to minimize a_p, because the cost of P is $1 - a_p + a_n$, which in limit is equal to $2 - a_p$, when n tends to infinity. We solved the resulting linear program and obtained a lower bound for POPoA converging to 1.99992 for $l = 1000$ and n tending to infinity. Thus, we have the following proposition:

Proposition 5.2 *There are games that have POPoA* 1.99992.

Here we note that one can obtain lower bound solving linear program directly, without introducing n which tends to infinity, but the speed of increasing lower bound is too low. For numerical comparison, the lower bound obtained for $n = 10{,}000$ is only 1.73. We leave it as an open problem to analyze the convergence of the POPoA of the above construction, and conjecture that it converges to two.

Conjecture 5.3 *There are games that have POPoA arbitrarily close to* 2.

6 Myopic sequential prices of anarchy/stability

In this section we study the myopic sequential price of anarchy and the myopic sequential price of stability.

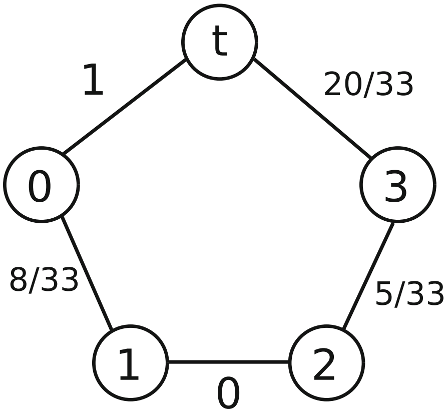

Fig. 13 Lower bound example for the sequential price of anarchy

6.1 Sequential price of anarchy in multicast games on rings

Lemma 6.1 *The myopic sequential price of anarchy is at most* 2 *in the multicast games on rings.*

Proof Consider an optimal strategy profile and let o be the edge that is not used. Consider any permutation (order) π of the players. If any player $\pi(i)$, $i < o$, decides to take a path containing edge o for the first time then it means that $a_o \leq \sum_{l=0}^{i} a_l$ which is bounded by the cost of the optimum. Therefore, the whole cost of the ring is bounded by 2 times the cost of the optimum. □

The presented upper bound is tight, as shows the example in Fig. 13, where $\pi = \{0, 1, 3, 2\}$ results in myopic sequential price of anarchy equal to 2.

6.2 Myopic sequential price of stability in multicast game

In the myopic sequential price of stability we consider the best permutation of players, with respect to the resulting network cost. In Bilò et al. (2010) authors prove that when the social knowledge network graph is directed acyclic then the price of anarchy is bounded by 4 (Theorem 8). If we consider that in the social knowledge graph each incoming player knows all the previous players then the result can be directly translated into our setting, but we give a different [simpler than the proof of general result in Bilò et al. (2010)] proof for our setting:

Theorem 6.2 *The myopic sequential price of stability in multicast games on arbitrary graphs is at most* 4.

Proof Since there is a common target vertex t, any optimum strategy profile forms a Steiner tree T on terminals s_i, $i = 0, \ldots, n-1$ and t. Consider a permutation of the vertices that corresponds to a depth-first search of the tree T, and make it the identity permutation $(0, 1, \ldots n-1)$. Let the players enter the game in this order, and make the myopic best responses. Denote by B_i the cost of the edges that player i uses alone in her strategy at the moment she enters the game, and let S_i be the overall cost of player i when she enters. Then the cost of the resulting network is $\sum_{i=0}^{n-1} B_i$. Since every player optimizes her cost when she enters the game, we have the following chain of inequalities: $S_i \leq d_T(s_i, s_{i-1}) + S_{i-1} - \frac{1}{2}B_{i-1}$, for $i = 1, \ldots, n-1$, where $d_T(u, v)$ is the distance between nodes u and v using only the edges of the tree T. Each player i has the following alternative strategy: first travel to the source (vertex) s_{i-1} using the edges of T, and then follow the strategy of player $i-1$. Note that in this alternative strategy, player i saves at least half of the cost of the edges that player $i-1$ takes alone when she enters the game. For the first player, we have the following inequality $S_0 \leq d_T(s_0, t)$, because when she enters the game, one of the possible strategies is to take a direct path from s_0 to t using only the edges of T. By summing up all inequalities given above, we get that $\frac{1}{2}\sum_{i=0}^{n-2} B_i + S_{n-1} \leq 2 \cdot cost(T)$. Note that $S_{n-1} \geq \frac{1}{2}B_{n-1}$, which results into the upper bound of 4. □

This upper bound is tight, as the example (Theorem 5) from Bilò et al. (2010) shows, ratio in the lower bound example is arbitrarily close to 4.

Proposition 6.3 *There is a multicast game with the myopic sequential price of anarchy arbitrarily close to* 4.

6.3 Myopic sequential price of stability on rings

In this section we consider the myopic sequential price of stability of the multicast games on rings. The example from Fig. 12 shows that it can be as high as $\frac{4}{3}$. We prove the following upper bound.

Theorem 6.4 *The myopic sequential price of stability in the multicast games on rings is at most* $\frac{26}{19}$.

Proof Assume that the optimum strategy profile does not include the edge of cost a_o, and without loss of generality $\sum_{i=0}^{o-1} a_i \geq \sum_{i=o+1}^{n} a_i$. Consider the permutation $\pi = \{n-1, \ldots, o, 0, 1, \ldots, o-1\}$. First $n-o$ players clearly take the right path, by our assumption. Consider the remaining players. If there is no player which, upon arrival, prefers the right path over the left path, then only edges of an optimum strategy profile are included into the resulting network which means that the myopic sequential price of stability is 1. If the very first player 0 prefers the right path, then all other players necessarily prefer the right path as well, and the resulting network consists of all edges except for that of weight a_0. But then a_0 is at least as large as a_o, resulting again the myopic sequential price of stability equal to 1. Suppose that there exists i such that every player $l \leq i$ prefers to take the left path, and only the player (vertex) $i+1$ prefers to take the right path. This implies the following inequalities:

$$\sum_{k=0}^{i} \frac{a_k}{i-k+1} \leq \sum_{k=i+1}^{n} a_k, \text{ and} \tag{8}$$

$$\sum_{k=i+2}^{o} a_k \leq \sum_{k=0}^{i+1} \frac{a_k}{i+2-k}, \tag{9}$$

where the first inequality (8) indicates that the i-th player prefers the left path, and the second inequality (9) indicates that the $i+1$-th player prefers the right. Our goal is to investigate the maximum possible cost c of the resulting network, where $c = a_0 + \cdots + a_i + a_{i+2} + \cdots + a_n$. Take the first inequality (8) with weight $\frac{2}{19}$, the second inequality (9) with weight $\frac{24}{19}$, and the normalization equation $a_0 + \cdots + a_{o-1} + a_{o+1} + \cdots + a_n = 1$ with weight $\frac{26}{19}$. We obtain that the sum on the left hand side s satisfies $c \le s \le \frac{26}{19}$, which gives that $c \le \frac{26}{19} \approx 1.368$. □

The permutation from the proof of Theorem 6.4 cannot be used to provide a better bound, as there exists an example of a game, where the permutation results in a network of cost $\frac{26}{19}$ times larger than the cost of the optimum. The example consists of 3 players. Edges on the ring have weights $\frac{6}{19}$, $\frac{10}{19}$, $\frac{3}{19}$ and $\frac{10}{19}$ in the counter-clockwise order. Players who come in the game according to the permutation $\{0, 1, 2, 3\}$ take all edges except for the 3-rd edge of weight $\frac{3}{19}$, resulting into a network of cost $\frac{26}{19}$, while the optimum solution cost is 1. Note that if players come according to the "opposite" permutation $(n-1, \ldots, 0)$, then the resulting network has the same cost as the optimum network. We have experimentally checked these two permutations, and for all inputs we tried, one of the two permutations resulted in a network of cost no more than the 4/3 of the optimum cost. Actually, we have checked that there is no instance of at most 1000 players where the better of the two permutations fails in that respect.

Conjecture 6.5 *The myopic sequential price of stability in the multicast game on rings is at most* $\frac{4}{3}$.

7 Conclusions

We reduced upper bound of the price of stability in the general network design game by analyzing a general optimum solution structure.

We have analyzed several solution concepts for the multicast network design games on rings, and demonstrated that they differ in terms of quality. Some of the derived bounds are not shown to be tight, and we leave it for future work to make them tight.

We have also initiated the study of the myopic sequential price of stability, and analyzed it for the multicast network design game on a ring. It is certainly an interesting challenge to provide better bounds on this concept for general (not multicast) network design games.

Acknowledgements We are grateful to anonymous reviewers for their excellent comments that improved the write-up of this paper. In particular, we are grateful to one anonymous reviewer for detecting a mistake in the proof of theorem 5.1 and for suggesting an alternative solution. We are grateful to Rati Gelashvili for valuable discussions and remarks. We used the CGAL linear and quadratic programming solver (Fischer et al. 2017) for solving all linear programs described in this article. This work has been partially supported by the Swiss National Science Foundation (SNF) under the Grant Number 200021_143323/1.

A Weights for inequalities from the Proof of Theorem 4.2

In this appendix we provide the multiplicative weights of the inequalities using a dual to a linear program that was solved to upper bound the price of stability in the multicast game on rings. The first inequality is the normalization inequality, therefore its weight is the upper bound on the price of stability. The next k inequalities indicate that the first $k \leq 7$ players left of edge e prefer to deviate, i.e., prefer to choose the right path instead of the left path, and the last inequality indicates that we have a Nash equilibrium, i.e., the last player considered in the best-response dynamics prefers to stick with the left path than to switch to the right path. The objective of the linear program is to minimize the sum of the edge costs without the edge that is not used by the Nash equilibrium achieved via the best response dynamics. The coefficients (weights) are as follows:

- $k = 1$ (0: 4/3; 1: 10/9; 2: 2/9)
- $k = 2$ (0: 22/17; 1: 252/323; 2: 202/323; 3: 90/323)
- $k = 3$ (0: 29/23; 1: 2976/4025; 2: 1206/4025; 3: 2256/4025; 4: 1224/4025)
- $k = 4$ (0: 1.243533565; 1: 0.722076586; 2: 446,160/1,659,763; 3: 0.268809463; 4: 0.528169383; 5: 0.329251827)
- $k = 5$ (0: 1.229596836; 1: 0.711037768; 2: 0.257115234; 3: 0.201170436; 4: 0.199302216; 5: 0.50797093; 6: 0.348431623)
- $k = 6$ (0: 1.217310111; 1: 0.702648246; 2: 0.250967669; 3: 0.189905238; 4: 0.168566505; 5: 0.179311025; 6: 0.494134279; 7: 0.362553601)
- $k = 7$ (0: 1.206536915; 1: 0.69586637; 2: 0.247111078; 3: 0.184286036; 4: 0.157438535; 5: 0.148587957; 6: 0.165607593; 7: 0.484007846; 8: 0.3733 84452)

For $k > 7$, we take only the first 7 inequalities indicating that the first 7 players prefer to take the right path than to stick to the left path. This is enough to prove an upper bound of 1.33081 for the price of stability. In the following, we list the weights of the inequalities of the dual to our linear program (index k : denotes the weight of the inequality to player k): (0: 1.330802428; 1: 0.750587484; 2: 0.246845878; 3: 0.168106752; 4: 0.12615003; 5: 0.096800836; 6: 0.072578056; 7: 0.048719834).

References

Adamy U, Ambühl C, Anand RS, Erlebach T (2007) Call control in rings. Algorithmica 47(3):217–238

Alós-Ferrer C, Netzer N (2010) The logit-response dynamics. Games Econ Behav 68(2):413–427

Angelucci A, Bilo V, Flammini M, Moscardelli L (2015) On the sequential price of anarchy of isolation games. J Comb Optim 29(1):165–181

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: FOCS, pp 295–304

Asadpour A, Saberi A (2009) On the inefficiency ratio of stable equilibria in congestion games. In: WINE, pp 545–552

Attiya H, Snir M, Warmuth MK (1988) Computing on an anonymous ring. J ACM 35(4):845–875

Bilò V, Bove R (2011) Bounds on the price of stability of undirected network design games with three players. J Interconnect Netw 12(1–2):1–17

Bilò V, Caragiannis I, Fanelli A, Monaco G (2013) Improved lower bounds on the price of stability of undirected network design games. Theory Comput Syst 52(4):668–686

Bilò V, Fanelli A, Flammini M, Moscardelli L (2010) When ignorance helps: graphical multicast cost sharing games. Theor Comput Sci 411(3):660–671

Bilò V, Flammini M, Moscardelli L (2014) The price of stability for undirected broadcast network design with fair cost allocation is constant. Games Econ Behav. https://doi.org/10.1016/j.geb.2014.09.010

Charikar M, Karloff HJ, Mathieu C, Naor J, Saks ME (2008) Online multicast with egalitarian cost sharing. In: SPAA 2008: proceedings of the 20th annual ACM symposium on parallelism in algorithms and architectures, Munich, Germany, June 14–16, 2008, pp 70–76

Chekuri C, Chuzhoy J, Lewin-Eytan L, Naor J, Orda A (2006) Non-cooperative multicast and facility location games. In: Proceedings 7th ACM conference on electronic commerce (EC), Ann Arbor, Michigan, USA, June 11–15, pp 72–81

Christodoulou G, Chung C, Ligett K, Pyrga E, van Stee R (2009) On the price of stability for undirected network design. In: WAOA, pp 86–97

Disser Y, Feldmann AE, Klimm M, Mihalák M (2015) Improving the Hk-bound on the price of stability in undirected shapley network design games. Theor Comput Sci 562:557–564

Fanelli A, Leniowski D, Monaco G, Sankowski P (2015) The ring design game with fair cost allocation. Theor Comput Sci 562:90–100

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, pp 608–618

Fischer K, Gärtner B, Schönherr S, Wessendorp F (2017) Linear and quadratic programming solver. CGAL Editorial Board

Jian L (2009) An upper bound on the price of stability for undirected shapley network design games. Inf Process Lett 109:876–878

Kawase Y, Makino K (2013) Nash equilibria with minimum potential in undirected broadcast games. Theor Comput Sci 482:33–47

Lee E, Ligett K (2013) Improved bounds on the price of stability in network cost sharing games. In: EC, pp 607–620

Leme RP, Syrgkanis V, Tardos É (2012) The curse of simultaneity. In: Innovations in theoretical computer science 2012, Cambridge, MA, USA, January 8–10, 2012, pp 60–67

Monderer D, Shapley LS (1996) Potential games. Games Econ Behav 14(1):124–143

Rosenthal RW (1973) A class of games possessing pure-strategy nash equilibria. Int J Game Theory 2(1):65–67