

Do Capacity Constraints Constrain Coalitions?

MICHAL FELDMAN and OFIR GERI, Tel Aviv University

We study strong equilibria in symmetric capacitated cost-sharing connection games. In these games, a graph with designated source s and sink t is given, and each edge is associated with some cost. Each agent chooses strategically an s - t path, knowing that the cost of each edge is shared equally between all agents using it. Two settings of cost-sharing connection games have been previously studied: (i) games where *coalitions* can form, and (ii) games where edges are associated with *capacities*; both settings are inspired by real-life scenarios. In this work we combine these scenarios and analyze strong equilibria (profiles where no coalition can deviate) in capacitated games. This combination gives rise to new phenomena that do not occur in the previous settings. Our contribution is twofold. First, we provide a topological characterization of networks that always admit a strong equilibrium. Second, we establish tight bounds on the efficiency loss that may be incurred due to strategic behavior, as quantified by the strong price of anarchy (and stability) measures. Interestingly, our results qualitatively differ from those obtained in the analysis of each scenario alone, and the combination of coalitions and capacities entails the introduction of more refined topology classes than previously studied.

CCS Concepts: • **Theory of computation** → **Quality of equilibria**; **Network formation**; *Algorithmic game theory*; *Solution concepts in game theory*; *Network games*;

Additional Key Words and Phrases: Capacities, coalitions, cost-sharing games, network congestion games, network design games, network topology, strong equilibrium, strong price of anarchy

ACM Reference Format:

Michal Feldman and Ofir Geri. 2016. Do capacity constraints constrain coalitions? ACM Trans. Econ. Comput. 5, 1, Article 8 (October 2016), 26 pages.

DOI: <http://dx.doi.org/10.1145/2955090>

1. INTRODUCTION

The construction of networks by autonomous agents can be significantly affected by strategic behavior. These situations are frequently modeled as *cost-sharing* connection games, which have been extensively studied [Albers 2009; Anshelevich et al. 2008a, 2008b; Epstein et al. 2009; Feldman and Ron 2015]. For example, consider the construction of a large computer network used by different countries, where each country is an autonomous system that serves its own strategic interests.

In cost-sharing games, a network is given and each edge is associated with a cost. Each one of n agents wishes to construct a path between its source and sink nodes, where the cost of each edge is shared equally between the agents who use it; each agent

An earlier version of this work appeared in the *Proceedings of the 29th AAAI Conference on Artificial Intelligence* (AAAI-15). This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 337122. Authors' addresses: M. Feldman, Blavatnik School of Computer Science, Schreiber Building, Tel Aviv University, P.O.B. 39040, Tel Aviv 69978, Israel; email: michal.feldman@cs.tau.ac.il; O. Geri (current address), Computer Science Department, Stanford University, Gates Building, 353 Serra Mall, Stanford, CA 94305, USA; email: ofirgeri@cs.stanford.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 2167-8375/2016/10-ART8 \$15.00

DOI: <http://dx.doi.org/10.1145/2955090>

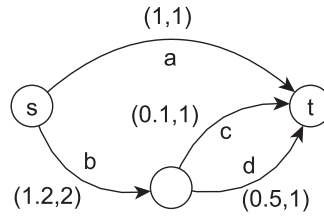


Fig. 1. A game with no SE.

desires to minimize his individual cost. Returning to our motivating example, a large computer network may be used by different countries, which should jointly cover its cost. All countries wish to use the network links but prefer to do so at minimal cost.

The analysis of these games evolved around the Nash equilibrium (NE) notion—a profile of strategies in which no agent can benefit from a unilateral deviation. In recent years, two interesting cases of these games have been considered, both inspired by real-life scenarios. First, a group of agents may form a *coalition* and collaborate for the benefit of all the members in the coalition. This scenario is formalized using the notion of a *strong equilibrium* (SE) [Aumann 1959]. A SE is a strategy profile from which no coalition can deviate in a way that benefits each one of its members. Note that we are considering non-cooperative games with non-transferable utilities. SE in cost-sharing connection games have been studied by Epstein et al. [2009]. The second scenario considers capacity constraints on the network edges [Feldman and Ron 2015]. Here, each edge is associated with some capacity that limits the number of agents who can use it. Both scenarios can naturally occur in our motivating example. Indeed, countries may collaborate to improve their standing, and capacity constraints may arise due to bandwidth limits.

While each of these settings has been previously studied alone, the combination of the two has not been previously considered. This is the focus of the present paper. Our main focus is on symmetric games, where all agents share the same source and sink nodes. Interestingly, the combination of coalitions and capacities gives rise to new phenomena that do not occur in any of the previous scenarios.

Equilibrium existence. In the setting with no capacities or coalitions, the profile in which all agents use the lowest-cost path from the source to the sink is clearly a NE. Now consider games with coalitions or capacities. The profile just mentioned is clearly a SE as well, as no coalition can benefit by deviating from the lowest-cost path. As for capacities, while capacity constraints may limit the use of the lowest-cost path, it has been proven that every (feasible) capacitated game admits a NE, due to the existence of an exact potential function [Feldman and Ron 2015]. Consider next the combination of coalitions and capacities. Here even the mere existence problem becomes non-trivial. Consider, for example, the game depicted in Figure 1, played by two agents. In our figures, we use the notation (x, y) to denote an edge with cost x and capacity y . In this game, the unique NE (up to renaming of agents) is one where one agent uses the upper edge (a) at cost 1, and the other agent uses the path b, c at cost 1.3. However, this profile is not a SE, as the two agents can deviate to the paths b, c and b, d , respectively, at the respective costs of 0.7 and 1.1. Since every SE is also a NE, and the only NE is not a SE, we conclude that this game does not admit a SE.

Efficiency loss. As another example, consider the efficiency loss incurred due to strategic behavior. The standard measure used to quantify the efficiency loss is the price of anarchy (PoA), defined as the ratio between the cost of the worst-case NE and the optimal cost (where the social cost is defined as the overall cost of the edges in use).

Table I. Upper Bounds for the PoA in Different Scenarios

		SPP and EP	SP	General
PoA	Uncap.	n	n	n
	Cap.	n	n	unbounded
SPoA	Uncap.	1	1	1
	Cap. (*)	H_n	n	unbounded
SPoS	Cap. (*)	H_n	$\Theta(n)$	unbounded

Cap. and uncap. are shorthands for capacitated and uncapacitated. Our results are marked with an asterisk (*). All bounds are tight.

In the game with no capacities or coalitions, it has been shown that the PoA can be as bad as n (i.e., the number of agents), and this is tight [Anshelevich et al. 2008b]. The consideration of coalitions significantly reduces the loss; in fact, every SE is optimal. In capacitated games, the PoA depends on the network topology. For example, if the network topology adheres to a *series-parallel* structure (defined in Section 2.2), then the PoA is bounded by n . Do coalitions reduce the efficiency loss in capacitated games as in their uncapacitated counterparts? Interestingly, the answer is *no* for series-parallel networks, but the answer is affirmative for a smaller set of topologies. Once again, the combination of coalitions and capacities introduces interesting phenomena and requires a more refined classification of network topologies than each setting alone.

1.1. Our Results

Equilibrium existence. As mentioned above, not all symmetric capacitated cost-sharing games admit a SE. We provide a full characterization of network topologies that do admit a SE; that is, every game played on a topology in this class admits a SE, and for every topology not in this class, there exists a game that does not possess a SE. The analysis of this part requires the introduction of a new class of networks, which we refer to as *Series of Parallel Paths* (SPP) networks. This class is defined as all networks that are the concatenation of parallel-path networks.

Efficiency loss. Previous analysis of the efficiency loss in cost-sharing games showed that the network topology significantly affects the incurred loss. This observation is reinforced in this work, as summarized below (see also Table I for a subset of our results). We provide tight bounds on the strong price of anarchy (SPoA), defined as the ratio between the cost of the worst-case SE and the optimal cost and on the strong price of stability (SPoS), defined analogously with respect to the best-case SE.

Epstein et al. [2009] establish an upper bound of H_n (i.e., the n th harmonic number, which is roughly $\log(n)$) on the SPoA in every game (including asymmetric games) that admits a SE. As mentioned above, this result does not carry over to capacitated networks. However, we show that this result does carry over to a specific class of network topologies, namely extension-parallel (EP) and SPP networks. Moreover, we provide an example showing that this bound is tight. In series-parallel (SP) networks, the SPoA is at most n (follows from the upper bound on the PoA), and we provide an example showing that this is tight. For general networks, we show that the SPoA can be arbitrarily high, even in instances with only two agents. Interestingly, our analysis results in a more refined classification of topologies than was required for the study of capacities or coalitions alone, most notably in the distinction between subclasses of SP networks.

In addition, we provide bounds on the SPoS. We show that the SPoS can also be as high as $\Omega(n)$ in SP networks and is unbounded in general networks. Interestingly, while the price of stability (PoS) is significantly better than the PoA (1 versus n in uncapacitated games, and an even wider gap in capacitated games), in capacitated

games we show that for all the topology classes we consider, the bounds on the SPoS and SPoA are asymptotically the same. A natural interpretation of the PoS measure is the loss that is incurred if there exists a coordinator who can suggest an initial configuration to the agents. While a coordinator can sometimes reduce the efficiency loss, our results here imply that a coordinator may not be useful in the worst case.

In addition to studying games with arbitrary capacity constraints, we consider the special case in which all the edges have the same capacity. We show that limiting the capacities to be homogeneous reduces the efficiency loss in games played on EP networks, in which every SE becomes a socially optimal solution.

Extensions. We consider two natural extensions to our model. First, we consider asymmetric games. Here, we provide a characterization of network topologies that always admit a SE and provide an example of a simple EP network for which the SPoA is unbounded. Second, we show that all of our results regarding symmetric games extend to undirected networks as well.

1.2. Related Work

Since their introduction by Anshelevich et al. [2008b], cost-sharing connection games have been widely studied in various settings. In the original setting that was studied by Anshelevich et al. [2008b], the agents' strategies consist of the amounts each of them is willing to pay for each edge, and in order to use an edge, the total amount that is paid for it should exceed its cost. This allows for general cost-sharing mechanisms that are not necessarily fair. The special case of fair cost-sharing was studied by Anshelevich et al. [2008a]. Under the fair cost-sharing rule, the strategy of an agent simply becomes a path in the graph, and the cost of each edge is shared equally between the agents who use it. They observed that the fair connection game belongs to the more general class of congestion games, which was introduced by Rosenthal [1973]. In these games, a pure NE always exists, and they are known to admit an exact potential function, as defined by Monderer and Shapley [1996].

The above works studied the inefficiency of equilibria by considering the measures of price of anarchy (PoA) and price of stability (PoS). The PoA measure was introduced by Koutsoupias and Papadimitriou [2009] and has since been studied in numerous games and settings. Anshelevich et al. [2008b] showed that the PoA in every cost-sharing connection game is at most n (the number of agents). It was also shown that the PoA can be as high as n even in a simple network that consists of two parallel edges. For fair cost-sharing games in directed networks, a tight upper bound of H_n on the PoS was established by Anshelevich et al. [2008a] using the potential function of the game. In the case of undirected networks, an upper bound of $O(\frac{\log n}{\log \log n})$ on the PoS in single-source games was established by Li [2009]. In the special case of broadcast games, in which there is a single source and each other node is the sink node of a different agent, the PoS was shown to be a constant [Bilò et al. 2013]. Lower bounds on the PoS in the undirected case were studied by Bilò et al. [2013], who showed that the PoS can be as high as $348/155$, 1.862 , and $20/11$ in general games, single-source games, and broadcast games, respectively.

The concept of a strong equilibrium—a strategy profile that is resilient to coalitional deviations—was introduced by Aumann [1959]. Epstein et al. [2009] studied SE in fair and general cost-sharing connection games. Their results showed that the network topology plays a major role in the existence of a SE. For example, they showed that every single-source game that is played on a series-parallel graph has a SE, but this property does not hold for single-source games played on general graphs. The existence of SE in the more general class of network congestion games was studied by Holzman and Law-Yone [1997, 2003] and by Holzman and Monderer [2015], who established

a characterization of the network topologies that admit SE in congestion games with monotone cost functions.

For fair cost-sharing, Epstein et al. [2009] and Albers [2009] independently proved an upper bound of H_n on the strong price of anarchy (SPoA), which is the analogue of the PoA with respect to SE. This upper bound holds for every game that has a SE, regardless of the network topology. The SPoA measure was introduced by Andelman et al. [2009] and was also studied with respect to various additional settings; see, for example, [Chien and Sinclair 2009] and [Fiat et al. 2007].

The consideration of edge capacities in cost-sharing games was first suggested by Feldman and Ron [2015]. They studied the PoA and PoS in symmetric games with regard to two objective functions: the sum of costs of all agents and the maximum cost of an agent. Their work has further emphasized the importance of the network topology. For example, when considering the sum of costs objective function, the PoA is at most n in games played on series-parallel networks but can be arbitrarily high in non-SP networks. A recent paper by Erlebach and Radoja [2015] showed that the PoS with respect to the max-cost objective function is at most n (in symmetric games), closing a gap that was presented by Feldman and Ron [2015].

2. MODEL AND PRELIMINARIES

2.1. Symmetric Capacitated Fair Cost-Sharing Connection Games

A symmetric capacitated fair cost-sharing connection (CF CSC) game is given by a tuple

$$\Delta = (n, G = (V, E), s, t, \{p_e\}_{e \in E}, \{c_e\}_{e \in E}),$$

where n is the number of agents, $G = (V, E)$ is a directed graph, $s, t \in V$ are the source and sink nodes (respectively), and each edge $e \in E$ is associated with a cost $p_e \in \mathbb{R}^{\geq 0}$ and a capacity constraint $c_e \in \mathbb{N} \cup \{0\}$. Each agent wishes to construct an s - t path in G while maintaining minimal cost. The strategy space of agent j , denoted by Σ_j , is the set of all s - t paths in G . The joint strategy space is denoted by $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$. The game is non-cooperative with non-transferable utilities.

Given a strategy profile $\mathbf{s} = (s_1, \dots, s_n) \in \Sigma$, the number of agents that use an edge e in the profile \mathbf{s} is denoted by $x_e(\mathbf{s}) = |\{j | e \in s_j\}|$. A profile \mathbf{s} is said to be feasible if for every $e \in E$, $x_e(\mathbf{s}) \leq c_e$. A game is said to be feasible if it admits a feasible strategy profile. Throughout this paper we only consider feasible games. When all the edges have the same capacity, we say that the game has homogeneous capacities. Formally, in a game with homogeneous capacities, there is a number $c \in \mathbb{N}$ such that for every $e \in E$, $c_e = c$.

For a given strategy profile \mathbf{s} and a coalition C , the induced strategy profile on the agents of the coalition C is denoted by \mathbf{s}_C , and the strategy profile of the rest of the agents is denoted by \mathbf{s}_{-C} . We consider the fair cost-sharing mechanism, where the cost of each edge is shared equally between all the agents who use it. The cost of agent j in a strategy profile \mathbf{s} is

$$p_j(\mathbf{s}) = \begin{cases} \sum_{e \in s_j} \frac{p_e}{x_e(\mathbf{s})} & \text{if } \mathbf{s} \text{ is feasible} \\ \infty & \text{otherwise} \end{cases}.$$

We use the utilitarian objective function, that is, the social cost of a strategy profile \mathbf{s} is the sum of costs of all agents, $cost(\mathbf{s}) = \sum_j p_j(\mathbf{s})$. The social cost of a profile is also equal to the sum of costs of the edges in use.

A strategy profile \mathbf{s} is a Nash equilibrium (NE) if no agent can improve her cost by deviating to another strategy, that is, for every j and every strategy $s'_j \in \Sigma_j$, it holds that $p_j(\mathbf{s}) \leq p_j(s'_j, \mathbf{s}_{-j})$ (where \mathbf{s}_{-j} denotes the strategy profile of all agents except j in \mathbf{s}). A strong equilibrium (SE) is a strategy profile in which no coalition can deviate

jointly in a way that will strictly decrease the cost of every coalition member. Formally, a profile \mathbf{s} is a SE if for every coalition of agents C and every set of strategies $\mathbf{s}'_C \in \Sigma_C$ there exists an agent $j \in C$ such that $p_j(\mathbf{s}) \leq p_j(\mathbf{s}'_C, \mathbf{s}_{-C})$. The sets of NE and SE of a game Δ are denoted by $NE(\Delta)$ and $SE(\Delta)$, respectively.

We use the price of anarchy (PoA) and price of stability (PoS) measures to quantify the efficiency loss incurred due to strategic behavior. Let \mathbf{s}^* be a strategy profile with minimal social cost in a game Δ . Then, the PoA of Δ is the ratio between the cost of the worst-case NE and the cost of \mathbf{s}^* , namely $PoA = \max_{\mathbf{s} \in NE(\Delta)} \frac{cost(\mathbf{s})}{cost(\mathbf{s}^*)}$. Similarly, the PoS is the ratio between the cost of the best-case NE and the cost of \mathbf{s}^* , namely $PoS = \min_{\mathbf{s} \in NE(\Delta)} \frac{cost(\mathbf{s})}{cost(\mathbf{s}^*)}$. The analogues of the PoA and PoS with respect to SE are named the strong price of anarchy (SPoA) and strong price of stability (SPoS). Formally, $SPoA = \max_{\mathbf{s} \in SE(\Delta)} \frac{cost(\mathbf{s})}{cost(\mathbf{s}^*)}$ and $SPoS = \min_{\mathbf{s} \in SE(\Delta)} \frac{cost(\mathbf{s})}{cost(\mathbf{s}^*)}$. For a family of games, these measures are defined with respect to the worst case over all the games in the family.

2.2. Graph Theoretic Preliminaries

A symmetric network is a graph $G = (V, E)$ with two designated nodes, a source $s \in V$ and a sink $t \in V$. Every node or edge in the network appears in at least one simple path from s to t . The networks we consider in this paper are directed unless stated otherwise. We hereby present three important operations on symmetric graphs.

- Identification: Given a graph $G = (V, E)$, the identification of two nodes $v_1, v_2 \in V$ yields a new graph $G' = (V', E')$, where $V' = (V \cup \{v\}) \setminus \{v_1, v_2\}$ and E' includes all the edges of E , where each edge that was connected to v_1 or v_2 is now connected to v instead. Figuratively, the identification operation is the collapse of two nodes into one.
- Series composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ with $s_1, t_1 \in V_1$ and $G_2 = (V_2, E_2)$ with $s_2, t_2 \in V_2$, the series composition $G = G_1 \rightarrow G_2$ is the network formed by identifying t_1 and s_2 in the union network $G' = (V_1 \cup V_2, E_1 \cup E_2)$. In the composed network G , the new source is s_1 and the new sink is t_2 .
- Parallel composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ with $s_1, t_1 \in V_1$ and $G_2 = (V_2, E_2)$ with $s_2, t_2 \in V_2$, the parallel composition $G = G_1 \parallel G_2$ is the network formed by identifying the nodes s_1 and s_2 (forming a new source s) and the nodes t_1 and t_2 (forming a new sink t) in the union network $G' = (V_1 \cup V_2, E_1 \cup E_2)$.

Using these operations, we define three classes of network topologies that will be of interest throughout the paper.

Definition 2.1. A symmetric network $G = (V, E)$ is a *series-parallel* (SP) network if it consists of a single edge or if there are two SP networks G_1, G_2 such that $G = G_1 \rightarrow G_2$ or $G = G_1 \parallel G_2$.

Definition 2.2. A symmetric network $G = (V, E)$ is an *extension-parallel* (EP) network if one of the following applies:

- (1) G consists of a single edge.
- (2) There are two EP networks G_1, G_2 such that $G = G_1 \parallel G_2$.
- (3) There is an EP network G_1 and an edge e such that $G = G_1 \rightarrow e$ or $G = e \rightarrow G_1$.

Definition 2.3. A symmetric network G is a *Series of Parallel Paths* (SPP) network if there exist networks G_1, \dots, G_k , each constructed by a parallel composition of simple paths, such that $G = G_1 \rightarrow \dots \rightarrow G_k$.

Note that every EP network is an SP network. An SPP network may or may not be an EP network, but it is always an SP network. Examples of SPP, EP, SP, and non-SP networks appear in Figures 11, 1, 6, and 2, respectively.

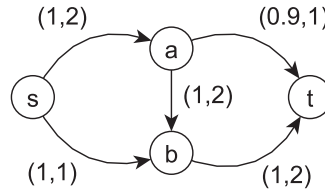


Fig. 2. A game with no SE played on a Braess graph.

Finally, we have to define when a network is embedded in another network. A symmetric network G is embedded in a network G' if G' is isomorphic to G or to a network derived from G using any number of the following operations:

- Subdivision: replacing an edge (u, v) by a new node w and two edges (u, w) and (w, v)
- Addition: adding a new edge connecting two existing nodes (including nodes that were added using subdivision or extension)
- Extension: adding a new source or sink node and an edge connecting the new node with the original source or sink node, respectively

3. EXISTENCE OF STRONG EQUILIBRIA

Every symmetric cost-sharing game admits a SE, as all agents can share the lowest-cost s - t path [Epstein et al. 2009]. In the capacitated version, it has been shown by Feldman and Ron [2015] that a pure NE exists in every feasible game by establishing that the game admits a potential function. Therefore, the consideration of capacities or coalitions alone does not preclude the existence of an equilibrium. However, as was already observed in the introduction (recall Figure 1), capacitated games may not admit any SE, even when played on a simple EP network (as shown in the example). We present an additional example of a game that does not admit a SE, which is played on an underlying network known as a Braess graph.

Example 3.1. Consider the game of two agents that is depicted in Figure 2. We show that each of the possible strategy profiles is not a SE. If both agents use the path (s, a, b, t) , one of the agents can profit by deviating to (s, a, t) . If one agent uses (s, a, t) and the other uses (s, a, b, t) , it is profitable for the latter to deviate to (s, b, t) . If one agent uses (s, a, b, t) and the other uses (s, b, t) , the former can reduce her cost by using (s, a, t) instead. A strategy profile in which one agent uses (s, a, t) and the other uses (s, b, t) is a pure NE, but it is not a SE because both agents can reduce their costs if they both use the path (s, a, b, t) . We conclude that this game does not admit a SE.

As it turns out, the two networks depicted in Figures 1 and 2 are, roughly speaking, the only barriers to SE existence. This is formalized in the remainder of this section as an exact characterization of the network topologies that always admit a SE.

Definition 3.2. A symmetric network G is said to admit a SE if every symmetric CFCS game played on G admits a SE.

Recall that a symmetric network is defined as a graph with designated source and sink nodes. According to the last definition, a network G is said to admit a SE if every CFCS game played on G (i.e., for every assignment of costs and capacities to the edges and for every number of agents) admits a SE. Conversely, a network G does not admit a SE if there exists an example of a game played on G that does not admit a SE. The following theorem establishes the characterization of networks that admit a SE.

ALGORITHM 1: Compute a SE for a Network of Parallel Edges.

-
- (1) $j \leftarrow 1$
 - (2) While there are agents that have not been assigned edges:
 - (a) For every edge e that has not been assigned to agents yet, compute its *fractional cost* $\frac{p_e}{\min\{c_e, n - \sum_{i=1}^{j-1} n_i\}}$.
 - (b) Find the edge e with the minimal fractional cost, and compute $n_j = \min\{c_e, n - \sum_{i=1}^{j-1} n_i\}$.
 - (c) Assign e to the following n_j agents: $\sum_{i=1}^{j-1} n_i + 1, \dots, \sum_{i=1}^j n_i$.
 - (d) Increment j .
-

THEOREM 3.3. *A symmetric network G admits a SE if and only if G is an SPP network.*

The proof is divided into two parts. In Theorem 3.4 we prove that every SPP network admits a SE, and in Theorem 3.6 we establish the converse direction.

THEOREM 3.4. *Every SPP network admits a SE.*

PROOF. We first observe that in our context, networks of parallel paths can be reduced to networks of parallel edges, where each path is replaced by an edge with capacity that equals the minimal capacity on the path and cost that equals the total cost of the path. Therefore, it suffices to prove the assertion of the theorem for networks constructed by series composition of parallel-edge graphs.

Our main building block in this proof is Algorithm 1, which computes a SE for games played on parallel-edge networks. The algorithm assigns edges to agents in a greedy manner. It first finds the edge with the lowest *fractional cost* (as defined in the description of the algorithm). When used to its maximum capacity (or by all the agents), this edge has the lowest possible cost for a single agent. The algorithm then assigns this edge to as many agents as possible. In the description of the algorithm, the number of agents who use the first edge is denoted by n_1 . In each iteration, the algorithm proceeds in the same way: It finds the edge with the lowest fractional cost that is still available and assigns it to as many agents as possible (this number is limited by the capacity of the edge and the number of remaining agents). The number of agents who use the j th edge is denoted by n_j . The algorithm terminates within at most n steps.

As part of the proof, we show that the returned profile is a SE. An important property of this algorithm is the *entireness property*: Each edge that is used by the algorithm is used to its maximum capacity, except for the last edge that is chosen by the algorithm, which may or may not be entirely used depending on the number of remaining agents. Moreover, the cost incurred due to an edge is exactly its fractional cost in the iteration in which it was chosen, and the costs of the agents are non-decreasing in the agents' indices.

After defining Algorithm 1, we are ready to prove the theorem. Let G_1, \dots, G_k be parallel-edge networks, and let G be a series composition of these networks. We define a strategy profile \mathbf{s} for the game played on G in the following way: For every $i = 1, \dots, k$, we compute a profile \mathbf{s}^i for network G_i using Algorithm 1. \mathbf{s} is defined to be the strategy profile in which the strategy of agent j is the concatenation of his strategies in $\mathbf{s}^1, \dots, \mathbf{s}^k$. We prove that \mathbf{s} is a SE.

From the definition of Algorithm 1, we get that for every profile \mathbf{s}^i , and for every two agents j, j' such that $j < j'$, $p_j(\mathbf{s}^i) \leq p_{j'}(\mathbf{s}^i)$. In other words, the first agent in \mathbf{s} incurs the lowest cost in each of G_1, \dots, G_k , the next agent incurs the second lowest cost in each of the subnetworks, and so on.

The intuition behind the proof is the following: Since the first agent pays the lowest possible cost in each of the subnetworks, she cannot reduce her cost in any subnetwork and thus cannot participate in any profitable coalitional deviation. Assuming that the first agent does not participate in the deviation, the second agent cannot reduce her cost in any subnetwork and cannot participate in the deviation as well. Arguments of this nature will be used to show that none of the agents can participate in the coalitional deviation.

Formally, assume by contradiction that there is a coalition C for which there is a profitable deviation, yielding a strategy profile \mathbf{s}' . Let j be the minimal index of an agent in C , that is, j is the agent that uses the j th lowest-cost edge in each of $\mathbf{s}^1, \dots, \mathbf{s}^k$.

In order for \mathbf{s}' to be a profitable coalitional deviation, j must reduce her cost in at least one of G_1, \dots, G_k . Assume, without loss of generality, that after the coalition deviates to \mathbf{s}' , agent j reduces her cost in G_1 . Denote the edge used by j in \mathbf{s}^1 by e . It follows that there exists an edge e' in G_1 such that

$$\frac{p_{e'}}{x_{e'}(\mathbf{s}')} < \frac{p_e}{x_e(\mathbf{s})}.$$

First, we show that $e' \neq e$. From the last inequality, if $e' = e$, the number of agents who use e in \mathbf{s}' must be larger than the number of agents who use e in \mathbf{s} . We consider each of the edges that are used in \mathbf{s}^1 and show that the number of agents who use it cannot increase after the deviation. If e is used to its maximum capacity, the number of agents who use it cannot increase. By the entirety property of Algorithm 1, we know that the only edge in \mathbf{s}^1 that may not be entirely used is the last edge that was picked by the algorithm. So if the edge e that agent j uses in \mathbf{s}^1 is the last edge chosen by the algorithm, all the agents with higher indices must also use e in \mathbf{s}^1 . Particularly, since j is the agent with the minimal index in C , it follows that all the agents in the coalition use e in \mathbf{s}^1 . Therefore, the number of agents who use e cannot increase after the deviation in this case as well. We have shown that the number of agents who use e in \mathbf{s}' cannot be larger than the number of agents who use e in \mathbf{s} and conclude that $e' \neq e$.

Second, we show that e' cannot be one of the edges that are entirely used in \mathbf{s}^1 . Since $e' \neq e$, e' was picked by the algorithm either before or after the iteration in which j was assigned an edge. Algorithm 1 picks edges in non-decreasing order of fractional cost, so agent j can only reduce her cost if e' was picked by the algorithm before j was assigned an edge. In that case, e' is only used by agents with indices lower than j . However, these agents do not participate in the deviation (j is the agent with the minimal index in C), and j cannot deviate to e' , which is used to its maximum capacity. Therefore, e' cannot be one of the edges that are entirely used in \mathbf{s}^1 .

We are left with two options for edge e' :

- (1) e' is partially used in \mathbf{s}^1 (i.e., e' is the last edge that was chosen by Algorithm 1). Since $e' \neq e$, the edge that agent j uses in \mathbf{s}^1 was picked by the algorithm before e' . So e' is used in \mathbf{s}^1 only by agents indexed $j + 1$ or higher. The only agents who may use e' in \mathbf{s}' are agents who use e' in \mathbf{s}^1 or agents that belong to C . These two sets of agents only contain agents with indices j or higher, thus $x_{e'}(\mathbf{s}') \leq |\{j, \dots, n\}| = n - j + 1$.
- (2) e' is not used in \mathbf{s}^1 . In this case, the agents who use e' in \mathbf{s}' must be members of the coalition $C \subseteq \{j, \dots, n\}$. Therefore, $x_{e'}(\mathbf{s}') \leq |\{j, \dots, n\}| = n - j + 1$.

In both cases, we have shown that $x_{e'}(\mathbf{s}') \leq n - j + 1$, and since the number of agents who use an edge is at most its capacity, $x_{e'}(\mathbf{s}') \leq \min\{c_{e'}, n - j + 1\}$. From the assumption that agent j reduces her cost in G_1 , we get that

$$\frac{p_{e'}}{\min\{c_{e'}, n - j + 1\}} < \frac{p_{e'}}{x_{e'}(\mathbf{s}')} < \frac{p_e}{x_e(\mathbf{s})},$$

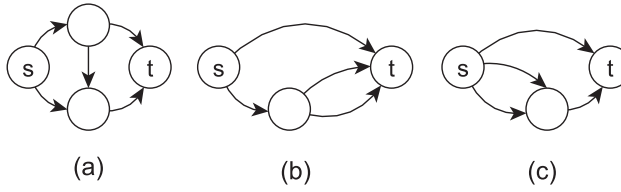


Fig. 3. Minimal non-SPP networks.

and the algorithm should have assigned e' to agent j (or a previous agent) instead of e . We get a contradiction and conclude that s is a SE. \square

We now prove the converse direction, namely that for every non-SPP network, there is a CFCSC game played on it that does not admit a SE. The key lemma in our proof connects us back to the two examples mentioned in the beginning of this section.

LEMMA 3.5. *A symmetric network is SPP if and only if it does not embed any of the networks depicted in Figure 3.*

PROOF. We first show that every SPP network does not embed any of the networks of Figure 3. Note that in a network of parallel paths, no node except for the source and the sink can have two edges entering or leaving it. Thus, a network of parallel paths cannot embed any of the networks of Figure 3. For an SPP network to embed one of the forbidden networks (the networks that appear in the figure), at least one of its subnetworks of parallel paths has to embed a forbidden network, which is impossible. We conclude that if a network is SPP, it cannot embed any of these networks.

In the converse direction, let G be a network that does not embed any of the networks of Figure 3. We prove that G is SPP. Holzman and Monderer [2015] proved that a network is SP if and only if it does not embed a Braess graph (Figure 3(a)). Therefore, G is SP. The rest of the proof is by induction on the height of the construction tree of G . The base case is a single edge, which is an SPP network. The induction hypothesis states that every SP network that does not embed any of the forbidden networks and has a construction tree with height smaller than that of G is SPP.

Assume, by contradiction, that G is not SPP. Since G is SP, we know that either $G = G_1 \rightarrow G_2$ or $G = G_1 \parallel G_2$. By induction, we can assume that both G_1 and G_2 are SPP (otherwise, one of G_1 and G_2 is an SP network with smaller construction tree that does not embed any of the forbidden networks and is not an SPP network). The series composition of two SPP networks is an SPP network, so it is impossible that $G = G_1 \rightarrow G_2$. Hence, $G = G_1 \parallel G_2$. If both of G_1, G_2 are networks of parallel paths, their parallel composition is a network of parallel paths, which is SPP. Therefore, one of G_1, G_2 embeds a series composition of a single edge and a parallel-edge network. It follows that G embeds one of the networks presented in Figures 3(b) and (c). It is a contradiction, hence G is SPP. \square

Using Lemma 3.5, we can now prove the following. The proof is similar to that of a lemma proven by Feldman and Ron [2015] and is specified here for completeness.

THEOREM 3.6. *For every non-SPP network G , there exists a symmetric CFCSC game played on G that does not admit a SE.*

PROOF. By Lemma 3.5, if G is not an SPP network, it must embed one of the networks depicted in Figure 3. We define a game of two agents played on G . We start from one of the networks presented in Figure 3 and assign costs and capacities as in Figures 1 and 2. The network G is derived from the forbidden network using the operations: subdivision, addition, and extension. We show how to assign costs and capacities to the

edges that are added in the process such that each strategy profile in the new game corresponds to a strategy profile in the original game, which does not admit a SE.

Formally, let Δ be a CFCSC game played on a network G , and let G' be a network derived from G by applying one of the operations: subdivision, addition, or extension. We construct a game Δ' played on G' that emulates Δ : We define a bijection f from the strategy space of an agent in Δ to the strategy space of an agent in Δ' , such that (s_1, \dots, s_n) is feasible whenever $(f(s_1), \dots, f(s_n))$ is feasible, and $p_j(s_1, \dots, s_n) = p_j(f(s_1), \dots, f(s_n))$ for every strategy profile (s_1, \dots, s_n) . We do not need to define different mappings between strategy spaces for different agents, as we consider symmetric games and the strategy spaces of the agents are identical. We define Δ' and the function f for each one of the possible operations:

- Subdivision: If an edge e is subdivided into e_1 and e_2 , we set $c_{e_1} = c_{e_2} = c_e$, $p_{e_1} = p_e$, and $p_{e_2} = 0$. Given a strategy s_j in Δ , $f(s_j)$ is the same path as s_j where the edge e is replaced by the path (e_1, e_2) (if $e \notin s_j$, $f(s_j) = s_j$). No strategy profile in Δ' can use only one of e_1, e_2 , so f is a bijection.
- Addition: If an edge e connecting two existing nodes is added, we set its capacity to be $c_e = 0$, and define $f(s_j) = s_j$ for every strategy s_j . No strategy in Δ' can use the new edge, so f is a bijection.
- Extension: If the source or the sink are extended using a new edge e , we set $c_e = n$ and $p_e = 0$. Given a strategy s_j , $f(s_j)$ uses the newly added edge e concatenated to the path used in s_j . Every strategy in Δ' must use the new edge e , concatenated to a path in Δ , so f is a bijection.

We claim that Δ' does not admit a SE if and only if Δ does not admit a SE. Let \mathbf{s}, \mathbf{s}' be strategy profiles in Δ , and let C be a coalition. The profile (s'_1, \dots, s'_n) is a profitable deviation from (s_1, \dots, s_n) for C in Δ , if and only if the profile $(f(s'_1), \dots, f(s'_n))$ is a profitable deviation from $(f(s_1), \dots, f(s_n))$ for C in Δ' . \square

This concludes the proof of Theorem 3.3.

4. STRONG PRICE OF ANARCHY

4.1. EP and SPP Networks

In this section, we bound the SPoA in capacitated games that admit SE. The following theorem establishes an upper bound on the SPoA for EP networks.

THEOREM 4.1. *For every symmetric CFCSC game played on an EP network, it holds that $SPoA \leq H_n$ (if a SE exists), and this bound is tight.*

PROOF. We first prove the upper bound. Epstein et al. [2009] showed that $SPoA \leq H_n$ for uncapacitated cost-sharing games. In their proof, they used the fact that no coalition can beneficially deviate from a strategy profile \mathbf{s} that is a SE (by definition). In particular, if some coalition C deviates to its corresponding profile in the socially optimal profile \mathbf{s}^* , then one of the agents in C weakly prefers the initial profile \mathbf{s} to the new profile $(\mathbf{s}_C^*, \mathbf{s}_{-C})$. The desired bound is then derived by the obtained inequalities for coalitions of sizes $n, \dots, 1$. The only barrier to applying the exact same technique to capacitated games is the fact that the deviation into profile $(\mathbf{s}_C^*, \mathbf{s}_{-C})$ might be infeasible due to capacity constraints. Our key lemma in this section shows that for games played on EP networks there always exists such a feasible deviation. The lemma is sufficient in order to prove the upper bound on the SPoA. Note that there may be more than a single socially optimal profile, but we have the freedom to choose one of them to be \mathbf{s}^* .

LEMMA 4.2. *Let G be an EP network and \mathbf{s} be a SE in a symmetric CFCSC game played on G . There exists a feasible strategy profile \mathbf{s}^* such that the cost of \mathbf{s}^* is minimal, and for every coalition C , the profile $(\mathbf{s}_C, \mathbf{s}_{-C}^*)$ is feasible.*

The following lemma, which is due to Feldman and Ron [2015], will be used in the proof.

LEMMA 4.3. [Feldman and Ron 2015] *Let G be an SP network. Let \mathbf{s} be a feasible profile of k agents in a game played on G , and let \mathbf{s}' be a feasible profile of r agents such that $r < k$. There exists an $s - t$ path in G that is feasible together with the strategies of the r agents in \mathbf{s}' and uses only edges that are used in the profile \mathbf{s} .*

PROOF OF LEMMA 4.2. Let G_{OPT} be the subnetwork that contains only the edges that are used by an optimal strategy profile, and let N denote the set of agents. We first define a specific profile \mathbf{s}^* played on G_{OPT} and then prove that for every coalition C , the strategy profile $(\mathbf{s}_C, \mathbf{s}_{-C}^*)$ is feasible. Since G is EP, G_{OPT} is also EP.¹ We define \mathbf{s}^* in two steps: First, we assign a strategy to as many agents as possible using recursion (on the structure of G_{OPT}). Then, we use Lemma 4.3 to extend this set of strategies to a profile of all agents. The profile \mathbf{s}^* is defined using Algorithm 2, which chooses a specific profile from all the optimal strategy profiles. The algorithm gets as input the subnetwork G_{OPT} that is used by an optimal profile and the strategy profile \mathbf{s} , which is a SE.

ALGORITHM 2: Choosing the Optimal Profile \mathbf{s}^*

Input: G_{OPT} is a graph, N is a set of agents, and \mathbf{s} is a strategy profile.

ChooseOptimalProfile(G_{OPT} , N , \mathbf{s}):

- (1) If $G_{OPT} = e$, where e is a single edge, return a strategy profile in which the edge e is assigned to all the agents in N who use e in \mathbf{s} .
 - (2) If $G_{OPT} = G_1 \parallel G_2$:
 - (a) Let $N_i \subseteq N$ be the set of agents that use an edge of G_i ($i = 1, 2$).
 - (b) $\mathbf{s}^1 \leftarrow \text{ChooseOptimalProfile}(G_1, N_1, \mathbf{s})$
 - (c) $\mathbf{s}^2 \leftarrow \text{ChooseOptimalProfile}(G_2, N_2, \mathbf{s})$
 - (d) Return the union of the profiles: $(\mathbf{s}^1, \mathbf{s}^2)$.
 - (3) If $G_{OPT} = G_1 \rightarrow e$ or $G_{OPT} = e \rightarrow G_1$, where e is an edge:
 - (a) $\mathbf{s}^1 \leftarrow \text{ChooseOptimalProfile}(G_1, N, \mathbf{s})$
 - (b) Each agent that has a strategy in \mathbf{s}^1 will also use the edge e .
 - (c) For any other agent that uses e in \mathbf{s} , attempt to find an available path in G_1 . If found, assign it together with the edge e to the agent.
 - (d) Return the profile that was defined in the last three steps.
 - In case it is possible to represent G_{OPT} as both $G_1 \rightarrow e_1$ and $e_2 \rightarrow G_2$, choose the representation in which the edge e_i is used by the maximal set of agents in \mathbf{s} .
-

It is important to note that the algorithm might possibly define a strategy profile only for a subset of the agents. In step 2, $N_1 \cap N_2 = \emptyset$, but it is possible that $N_1 \cup N_2 \subset N$ (since G_{OPT} is a subnetwork of G). Therefore, by the end of step 2, it is possible that not all the agents in N are assigned a strategy. In step 3(c), it is also possible that not all the agents in N will be assigned a strategy (as an agent is assigned a strategy only if there is a path available in G_1). Thus, we have to extend this strategy profile to a feasible profile for all agents (using only edges from G_{OPT}). We do so by applying

¹This can be observed by constructing G_{OPT} using the same construction tree of G , excluding parallel composition operations of subnetworks that all their edges are not used in G_{OPT} .

Lemma 4.3 to the network G_{OPT} , where \mathbf{s} is any strategy profile that uses only edges from G_{OPT} , and \mathbf{s}' is the partial profile that was defined above. This provides a full definition of the profile \mathbf{s}^* . Example 4.4 that follows this proof illustrates the definition of \mathbf{s}^* .

We claim that the profile \mathbf{s}^* satisfies the capacity constraints. In step 1 of the algorithm, an edge is assigned to the same agents that use it in the feasible profile \mathbf{s} (due to the way N is split in step 2). In step 3, the edge e is used only by agents that use it in \mathbf{s} , and edges in G_1 are assigned in a way that satisfies the capacity constraints. Therefore, no edge exceeds its capacity.

Let C be a coalition of agents. We prove that the strategy profile $\mathbf{s}_{comb} = (\mathbf{s}_C, \mathbf{s}_{-C}^*)$ is feasible. Let e be an edge. Let M denote the set of agents that use e in \mathbf{s} , and let M^* denote the set of agents that use e in \mathbf{s}^* . The set of agents that use e in \mathbf{s}_{comb} is $(M \cap C) \cup (M^* \cap (N \setminus C))$. If e is used only in \mathbf{s} , it cannot exceed its capacity in \mathbf{s}_{comb} , since $M \cap C \subseteq M$, and the profile \mathbf{s} is feasible. The same applies to edges that are used only in \mathbf{s}^* . It remains to consider edges that are used in both \mathbf{s} and \mathbf{s}^* . There are two types of edges in EP networks: ones that are added to the graph through an extension of the source or sink and all other edges. The algorithm assigns edges of the former type to agents in step 3 and edges of the latter type in step 1.

If edge e is used in both \mathbf{s} and \mathbf{s}^* , there are two cases:

- (1) If e is assigned to agents in step 1 of the algorithm, the set N in this step contains all the agents that use e in \mathbf{s} , namely, $M \subseteq N$, thus, $M \subseteq M^*$ ($e \in s_j \Rightarrow e \in s_j^*$). The set of agents that use e in \mathbf{s}_{comb} is $(M \cap C) \cup (M^* \cap (N \setminus C)) \subseteq M^*$. Since \mathbf{s}^* is feasible, then e does not exceed its capacity in \mathbf{s}_{comb} .
- (2) If e is added to the network G_{OPT} using extension of the source or the sink, the algorithm assigns e to agents in step 3. There are two sub-cases.
 - (a) In step 3(c), each agent that uses e in \mathbf{s} was assigned a path in G_1 in \mathbf{s}^* . In that case, $M \subseteq M^*$. The set of agents that use e in \mathbf{s}_{comb} is $(M \cap C) \cup (M^* \cap (N \setminus C)) \subseteq M^*$. Since \mathbf{s}^* is feasible, we get that e does not exceed its capacity in \mathbf{s}_{comb} .
 - (b) There is an agent that uses e in \mathbf{s} but does not use e and G_1 in \mathbf{s}^* . Due to the definition of \mathbf{s}^* , there is no available path in G_1 . Therefore, no agent can be assigned a path in G_1 later. Each agent that uses e in \mathbf{s}^* must use G_1 . Hence, no agents will be assigned the edge e after it is used in step 3 of the algorithm. Thus, $M^* \subseteq M$. The set of agents that use e in \mathbf{s}_{comb} is $(M \cap C) \cup (M^* \cap (N \setminus C)) \subseteq M$. Since \mathbf{s} is feasible, e does not exceed its capacity in \mathbf{s}_{comb} .

We conclude that no edge exceeds its capacity in \mathbf{s}_{comb} . \square

We turn to discuss the tightness of the upper bound established by the previous lemma. Feldman and Ron [2015] show an example of a game with n agents, in which the PoS is $\frac{H_n}{1+\epsilon}$ for every $\epsilon > 0$. Their example is included here as Figure 4. The only NE in this example is also a SE. This shows that the upper bound of H_n on the SPoA in EP networks is tight. \square

Example 4.4. Figure 5 provides an example that illustrates the definition of \mathbf{s}^* . Recall that Algorithm 2 gets as input a SE (\mathbf{s}) and the subnetwork that is used by an optimal strategy profile (G_{OPT}). The algorithm is used to choose an optimal profile \mathbf{s}^* from all possible optimal strategy profiles that use only edges from G_{OPT} . Figure 5(a) shows the strategies of $n = 6$ agents in the profile \mathbf{s} , which is a SE. Figures 5(b)–(d) show the construction of \mathbf{s}^* over G_{OPT} . Figure 5(b) illustrates the case of parallel composition: The lower edge is used by agent 6 (as in profile \mathbf{s}), and we use the algorithm recursively in order to find a profile for agents 1, \dots , 5 who use the upper subnetwork. Figure 5(c) shows the case of edge extension: The edge of cost 8 is a subnetwork that was extended

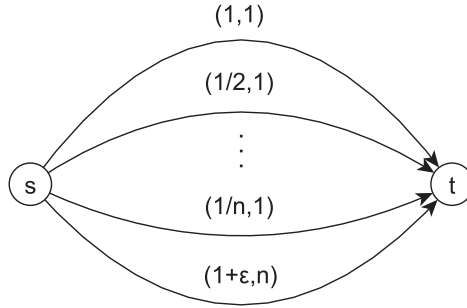
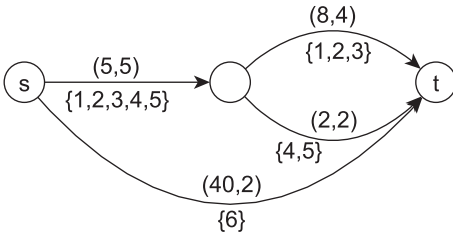
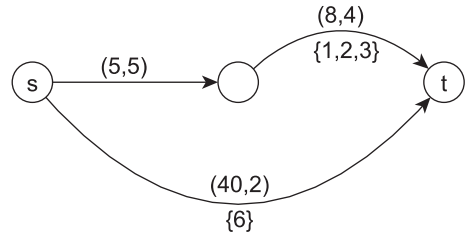


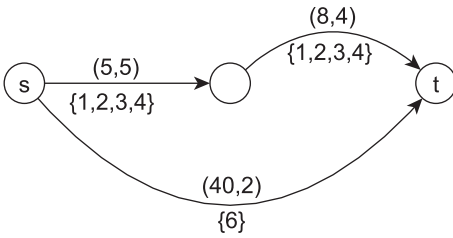
Fig. 4. [Feldman and Ron 2015] A game played on an EP network with $PoS = \frac{H_n}{1+\epsilon}$.



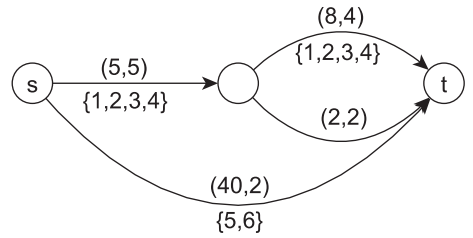
(a) The profile \mathbf{s} (a strong equilibrium)



(b) Cases 1 (a single edge) and 2 (parallel composition)



(c) Case 3 (edge extension)



(d) The final optimal profile \mathbf{s}^*

Fig. 5. An illustration of Algorithm 2.

with the edge of cost 5. The edge of cost 8 is used by the agents 1, 2, 3 in \mathbf{s} , and hence it will be used by the same agents in \mathbf{s}^* . The algorithm now defines which agents will use the edge of cost 5. First, it has to be used by agents 1, 2, 3. We consider the rest of the agents that use this edge in \mathbf{s} . One of them (agent 4) can use the edge of cost 8, and the other (agent 5) cannot. Therefore, agent 4 is assigned the same path as agents 1, 2, 3, and the strategy of agent 5 is left undefined. Figure 5(d) shows the final profile \mathbf{s}^* , after it was extended to all agents using Lemma 4.3.

Remark 4.5. Let $c_{max} = \max_{e \in E} c_e$ be the maximal capacity of an edge. In the case of capacitated games, the bound of H_n on the SPoA that was established by Epstein et al. [2009] can be slightly improved to $H_{c_{max}}$. In their proof, Epstein et al. [2009] used the potential function to bound the cost of every SE. For a SE \mathbf{s} and an optimal solution \mathbf{s}^* , they have shown that $cost(\mathbf{s}) \leq \Phi(\mathbf{s}^*)$, where the potential function Φ is defined by $\Phi(\mathbf{s}) = \sum_{e \in E} p_e \cdot H_{x_e(\mathbf{s})}$ (where $x_e(\mathbf{s})$ is the number of agents who use edge e in \mathbf{s}). When

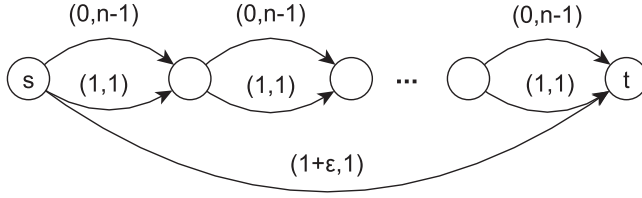


Fig. 6. A game played on an SP network with $SPoA \geq \frac{n}{1+\epsilon}$.

capacity constraints are added, in every feasible profile \mathbf{s} it holds that $x_e(\mathbf{s}) \leq c_{max}$. Therefore, we can bound the cost of every SE by $\Phi(\mathbf{s}^*) \leq \sum_{e \in E} p_e \cdot H_{c_{max}} = H_{c_{max}} \cdot cost(\mathbf{s}^*)$.

The following theorem extends the class of networks for which the SPoA is bounded by H_n .

THEOREM 4.6. *For every symmetric CFCSC game played on a network that is a series composition of EP networks, it holds that $SPoA \leq H_n$ (if a SE exists).*

PROOF. Let $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$ be a network. It is sufficient to prove that the SPoA of a game played on G is upper bounded by the maximal SPoA of a game played on one of the subnetworks G_1, \dots, G_k . Denote by $cost_{G_i}(\mathbf{s})$ the social cost of the strategy profile \mathbf{s} in a game played only on the subnetwork G_i , and let \mathbf{s}^* be an optimal profile. For every strategy profile \mathbf{s} , it holds that

$$cost_G(\mathbf{s}) = cost_{G_1}(\mathbf{s}) + \dots + cost_{G_k}(\mathbf{s}).$$

If \mathbf{s} is a SE in G , the induced strategy profile in G_i is also a SE. We get that

$$\begin{aligned} SPoA(G) &= \max_{\mathbf{s} \text{ is SE}} \frac{cost_G(\mathbf{s})}{cost_G(\mathbf{s}^*)} \\ &= \max_{\mathbf{s} \text{ is SE}} \frac{cost_{G_1}(\mathbf{s}) + \dots + cost_{G_k}(\mathbf{s})}{cost_{G_1}(\mathbf{s}^*) + \dots + cost_{G_k}(\mathbf{s}^*)} \\ &\leq \max_{\mathbf{s} \text{ is SE}} \max_i \frac{cost_{G_i}(\mathbf{s})}{cost_{G_i}(\mathbf{s}^*)} \\ &\leq \max_i SPoA(G_i). \quad \square \end{aligned}$$

Note that SPP networks are a special case of the class specified in the last theorem, and therefore the upper bound of H_n applies to SPP networks as well.

4.2. SP and General Networks

For SP networks, it was established by Feldman and Ron [2015] that $PoA \leq n$, which directly implies that $SPoA \leq n$. We provide an example showing that this bound is tight.

THEOREM 4.7. *For every $\epsilon > 0$, there exists an SP network G and a CFCSC game played on G such that $SPoA \geq \frac{n}{1+\epsilon}$.*

PROOF. Consider the game played by n agents on the graph depicted in Figure 6. The strategy profile in which each agent uses a different edge of cost 1 and $n - 1$ additional edges of cost 0 is a SE: In this profile, each agent pays 1, and the only way to reduce this cost is by using a path of edges that cost 0. However, the edges that cost 0 are already used to their maximum capacity, and any deviation that allows an agent to use a path that costs 0 will increase the cost of another agent.

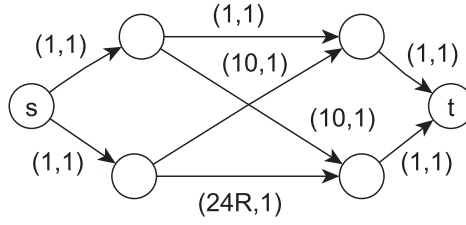


Fig. 7. A game with unbounded SPoA.

The optimal profile in this game is the profile where one agent uses the lower edge of cost $1 + \epsilon$ and $n - 1$ agents use the upper edges of cost 0. It follows that $SPoA \geq \frac{n}{1+\epsilon}$, as stated. \square

For general networks, the SPoA can be arbitrarily high, even with only two agents.

THEOREM 4.8. *For every real number R , there exists a CFCSC game with two agents in which the SPoA is greater than R .*

PROOF. Consider the game played by two agents on the graph presented in Figure 7. The strategy profile in which no agent uses the inner edges is a SE: The agent that uses the upper path incurs the lowest possible cost and will not participate in any deviation, while the other agent uses the only path that is still available. However, the optimal strategy profile will avoid the edge of cost $24R$. Thus, $SPoA \geq \frac{24R+5}{24} > R$. \square

4.3. Homogeneous Capacities

We consider a more restricted form of capacity constraints, in which all the edges have the same capacity. In this special case, we show that every SE in games played on EP networks is optimal.

THEOREM 4.9. *For every feasible symmetric CFCSC game with homogeneous capacities played on an EP network, it holds that $SPoA = 1$.*

PROOF. Let Δ be a feasible symmetric game played on an EP network G , and denote by c the capacity of the edges. We show that every SE in Δ is optimal. We first show that there is a game that is played on a network of parallel edges, such that every SE in that game is optimal if and only if every SE in Δ is optimal. Let G_1, \dots, G_k be EP networks such that $G = G_1 \parallel G_2 \parallel \dots \parallel G_k$, and none of G_1, \dots, G_k can be represented as a parallel composition of two networks. Since G_i is an EP network, and cannot be decomposed into two parallel subnetworks, G_i must be a single edge or an EP network that is extended by a single edge. It follows that in every strategy profile, no more than c agents can use edges that belong to G_i . If G_i is a single edge, its capacity is c , and if G_i is an EP network extended by a single edge e , every agent that uses edges from G_i must also use e , which has a capacity of c agents.

Next, we claim that in every SE, all the agents that use edges from G_i use the same path, which is the lowest-cost path in G_i (if there is more than one path with minimal cost, one of these paths will be shared by all the agents). If there are agents who use a path of higher cost in G_i , they can deviate to the lowest-cost path and reduce their costs. This deviation satisfies the capacity constraints, as at most c agents use edges from G_i , and the capacity of all the edges in the path is c .

Using the last observation, we can look at a different game that is played on a network of parallel edges and where every SE in Δ corresponds to a SE in the new game. For each of G_1, \dots, G_k , we compute the minimum cost of an s - t path that goes

through that subnetwork and replace the subnetwork with an edge that has that cost and a capacity of c agents. Given a SE in the original game, we construct a strategy profile for the new game, such that the agents who use the minimal cost path in G_i will use the edge that replaced G_i in the new network. The costs of these two strategy profiles are the same for each agent.

It is only left to show that when all the capacities are equal, every SE in a game played on a network of parallel edges is optimal. Let \mathbf{s} be a SE in a game played on a network of parallel edges, where the capacity of each edge is c . Denote by e the edge that has the maximal cost of all the edges that are used in \mathbf{s} . Assume (by contradiction) that there is another edge e' , such that $p_{e'} < p_e$ (e' is cheaper) and $x_{e'}(\mathbf{s}) < c$ (e' is available). The number of agents who can deviate from e to e' is at most $\min\{x_e(\mathbf{s}), c - x_{e'}(\mathbf{s})\}$. If this number of agents deviate to e' , their cost will reduce from $\frac{p_e}{x_e(\mathbf{s})}$ to at most $\frac{p_{e'}}{x_{e'}(\mathbf{s})}$. This contradicts our assumption that \mathbf{s} is a SE. Therefore, every SE uses $\lceil \frac{n}{c} \rceil$ edges of minimal total cost, as does the optimal solution. We conclude that $SPoA = 1$. \square

By applying the same arguments used in the proof of Theorem 4.6, the last result can be extended to every network that is a series composition of EP networks, particularly, SPP networks.

COROLLARY 4.10. *For every feasible symmetric CFCSC game with homogeneous capacities that is played on a network that is a series composition of EP networks, it holds that $SPoA = 1$.*

When considering more general networks, limiting the capacities to be homogeneous does not improve the SPoA. In the example presented in the proof of Theorem 4.7, which shows that the SPoA in SP networks can be arbitrarily close to n , we can replace each edge that costs 0 with $n - 1$ parallel edges that cost 0 and have unit capacity. This way, all the edges have unit capacity, and the theorem still holds for homogeneous capacities. In the example that shows that the SPoA can be unbounded (in non-SP networks), all the edges already have unit capacities.

5. STRONG PRICE OF STABILITY

In some cases, the best SE may be of interest as well, for example, if there exists a central entity that can coordinate the agents around an initial equilibrium. Clearly, it always holds that $PoS \leq SPoS \leq SPoA$ (since every SE is a NE), so the upper bounds on the SPoA apply to the SPoS as well. Interestingly, the upper bounds on the SPoS in all the classes of network topologies considered here match the upper bounds on the SPoA. This is in stark contrast to previous settings, which exhibited large gaps between worst-case and best-case equilibria. In what follows, we show that all the upper bounds on the SPoA established in the previous section are tight with respect to the best SE.

For EP and SPP networks, the bound of H_n is tight with respect to the best SE due to an example provided by Feldman and Ron [2015], in which $PoS = \frac{H_n}{1+\epsilon}$ for every $\epsilon > 0$. The example appears here in the proof of Theorem 4.1 (see Figure 4).

For SP networks, we have shown that $SPoA \leq n$, and the following theorem shows a game for which $SPoS = \Omega(n)$.

THEOREM 5.1. *There exists a symmetric CFCSC game played on an SP network, in which the SPoS is at least $\Omega(n)$.*

PROOF. Consider the game of n agents depicted in Figure 8. We divide the network into two parts. The first part, which we refer to as the upper part of the network, contains the paths that go from s to t through the nodes v_1, \dots, v_{n-2} . In the upper

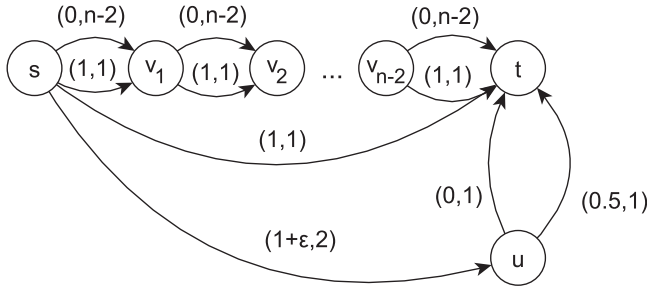


Fig. 8. A game played on an SP network with $SPoS = \Theta(n)$.

subgraph, there are $n - 1$ edges that cost 0 and $n - 1$ edges that cost 1. The second part, which we refer to as the lower part of the network, contains the edge connecting s and t directly and the two paths that go from s to t through node u . The upper part of the network can be used by at most $n - 1$ agents, so one of the agents must use the lower part of the network.

We show that the only SE in this game are profiles in which each agent pays 1: $n - 1$ agents use the upper part of the network. Each one of them uses a different edge of cost 1 and $n - 2$ edges of cost 0. The last agent uses the edge that goes directly from s to t and costs 1. We show that these profiles are SE by considering the possibly profitable deviations. There are two possible paths that cost less than 1. The first such path requires the agent to share the edge (s, u) with another agent, and then use the left edge (u, t) . However, a deviation of this kind would require another agent to use the path (s, u, t) that goes through the right edge (u, t) . That path costs $1 + \frac{\epsilon}{2}$, and the agent using it could not benefit from the deviation. The other path that costs less than 1 is the path that goes only through the edges that cost 0. However, these edges are used to their maximum capacity, and any coalitional deviation that allows an agent to use a path of cost 0 will not be profitable for one of the participating agents. Since there are no profitable deviations, the profiles are SE, and their social cost is n .

Now, we show that all other profiles are not SE. In every other profile in which only one agent uses the lower subgraph, one of the agents who use the upper part of the network pays at least 2. This agent can deviate to one of the lower paths and reduce her cost. Therefore, these profiles are not SE. Every profile in which there are two agents using the lower subgraph cannot be a SE either. In such profiles, the agents that use the upper part of the network will only use the edges that cost 0 and will not participate in any deviation. The remaining two agents always have a profitable deviation: In every NE, one of them will use the edge (s, t) that costs 1. However, this agent's cost can be reduced by sharing the edge (s, u) with the other agent, so that one of them pays $\frac{1+\epsilon}{2}$ and the other pays $1 + \frac{\epsilon}{2}$ instead of $1 + \epsilon$. It remains to prove that the profiles in which three agents use the lower subgraph are not SE, but these are not even NE, as one of the agents can benefit by deviating to a path consisting of edges of cost 0.

So far we have shown that the social cost of every SE is n . The optimal profile uses the edges that cost 0, and the two paths that go through node u . The social cost of this profile is $1.5 + \epsilon$, hence it holds that $SPoS = \frac{n}{1.5+\epsilon} = \Omega(n)$. \square

Next, we show that in the case of non-SP networks, the SPoS can be unbounded (as shown for the SPoA).

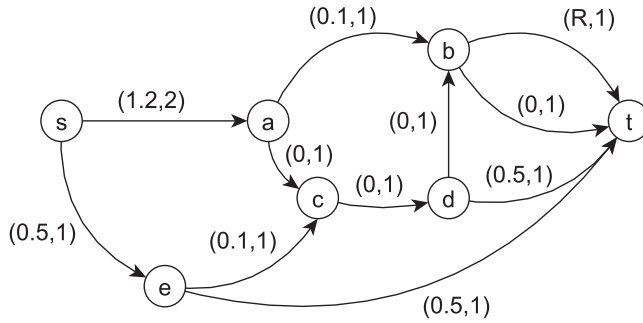


Fig. 9. A game with unbounded SPoS.

THEOREM 5.2. *There exists a symmetric CFCSC game with two agents in which the SPoS can be arbitrarily high.*

PROOF. Consider the game of two agents presented in Figure 9. The following strategy profile is a SE: One agent uses the path (s, e, c, d, b, t) and incurs the cost 0.6, and the other agent uses the path (s, a, b, t) and incurs the cost $R + 1.3$. This is a SE since the first agent incurs the minimal possible cost and the other agent uses the only available path. We claim that this is the only SE.

In every NE, it is impossible that two agents will use the edge (s, a) , as one of them incurs a cost of at least 1.1 and therefore may deviate to the path (s, e, t) . Hence, one agent must use the edge (s, a) and the other must use the edge (s, e) .

If one agent uses the path (s, e, t) , using the path (s, a, c, d, b, t) minimizes the cost of the other agent. This strategy profile is not a SE, as they can jointly deviate to the paths (s, a, b, t) and (s, a, c, d, t) , respectively, and both agents will benefit from the deviation.

We obtain that in every SE one agent must use the path (s, e, c, d) . This agent will not use the path (s, e, c, d, t) , as he can use the better path (s, e, t) . The agent will not use the edge that costs R for the same reason.

We conclude that the only SE is the one presented above. This SE uses the edge that costs R . An optimal strategy profile will avoid that edge, as R can be arbitrarily high. Therefore, the SPoS is unbounded, as stated. \square

6. EXTENSIONS

6.1. Asymmetric Games

A natural extension of our model is the case of asymmetric games, where different agents are associated with different source and sink nodes. In this section, we remark on the existence of SE and their quality in asymmetric games. Note that we still consider games that are played on networks with designated source and sink nodes (as defined earlier), but the source and sink nodes of network do not necessarily serve as the source and sink of the individual agents.

Regarding the existence of SE, Epstein et al. [2009] provided an example of an asymmetric game played on an SPP network that does not admit a SE. The following two theorems provide a characterization of the networks that admit a SE: one theorem considers games with a single source and the other considers games with multiple source nodes. The proofs are deferred to the appendix.

THEOREM 6.1. *In single-source games, a symmetric network G admits a SE if and only if it is SPP.*

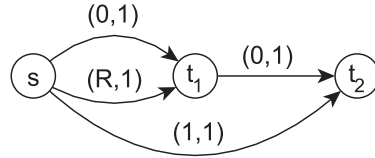


Fig. 10. An asymmetric game with unbounded SPoA.

THEOREM 6.2. *In games with multiple source and sink nodes, a symmetric network G admits a SE if and only if there are networks G_1, \dots, G_k such that $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$, and that one of G_1, \dots, G_k is a network of parallel paths and the rest are networks of parallel edges.*

Next we study the SPoA in asymmetric games. Unlike the symmetric case, which exhibited reasonable bounds for some families of topologies, in the asymmetric case, even if all agents share the same source (but not the same sink) and play a game on a simple EP network, the SPoA may be unbounded.

Example 6.3. Consider the game played by two agents that is presented in Figure 10. The source node of both agents is s , and the sink nodes are t_1 and t_2 , respectively. The strategy profile in which agent 1 uses the edge of cost R and agent 2 uses the path of cost 0 is a SE. In the optimal profile, agent 1 uses the edge of cost 0, and agent 2 uses the edge of cost 1. Therefore, in this example $SPoA = R$, which can be arbitrarily high.

However, for the class of SPP networks, the SPoA is upper bounded by H_n , even for asymmetric games.

THEOREM 6.4. *For every CFCSC game (either symmetric or asymmetric) played on an SPP network, it holds that $SPoA \leq H_n$ (if a SE exists).*

The proof is deferred to the appendix.

6.2. Undirected Graphs

Until now, we considered the case of cost-sharing games played on directed graphs. In this section, we show that our results (for symmetric games) extend to undirected graphs.

For directed graphs, we have shown that a network admits a SE if and only if it is an SPP network. Theorem 3.4 asserts that every SPP network admits a SE. To prove existence, we use a greedy algorithm for assigning edges to agents, such that in every coalition, there must exist an agent that cannot reduce his cost in any subnetwork of parallel edges. The same algorithm can be used to compute a SE in the undirected case (due to exactly the same arguments). In the converse direction, we presented two examples of games played on directed graphs that admit no SE. These games admit no SE also in the case where the underlying graph is undirected. To complete the argument, we note that the lemma that is used to extend these forbidden networks to networks that embed them applies to undirected graphs following an identical analysis.²

For bounding the SPoA and the SPoS in EP and SP networks, we use the following claim.

CLAIM 6.5. *Let Δ be a symmetric CFCSC game played on a directed SP network G , and let Δ' be the same game played on the same network but with undirected edges. If s is a SE in Δ' , then Δ has a SE with the same social cost as in s .*

²In the undirected case, we use the characterization of SP networks shown by Milchtaich [2006].

PROOF. Let \mathbf{s} be a SE in Δ' . We first claim that we can assume, without loss of generality, that \mathbf{s} uses only acyclic paths. To see this, note that the paths in \mathbf{s} cannot contain cycles that cost more than 0, as if one of them did, the agent using that path could reduce her cost by removing the cycle. In addition, it can be assumed that the paths in \mathbf{s} do not contain cycles that cost 0, as removing these cycles does not affect the agent costs. To conclude the proof, we use the following lemma, established by Milchtaich [2006].

LEMMA 6.6 [MILCHTAICH 2006]. *Let G be an undirected SP network, and let u, v be two distinct vertices in G . If there exists an acyclic s - t path in which u precedes v , then u precedes v in every acyclic s - t path that contains both vertices.*

By the above lemma, it follows that every edge used in the strategy profile \mathbf{s} (in the game Δ') is used in the direction determined by G . Each edge (u, v) in the directed SP network G can be used as part of an acyclic path in G , and the lemma guarantees that if u, v are both reached in a path in \mathbf{s} , u precedes v . Therefore, \mathbf{s} is a valid strategy profile in Δ . Clearly, any deviation that is feasible in Δ is also feasible in Δ' ; therefore the profile \mathbf{s} is a SE in Δ as well. \square

The above claim implies that the set of SE in the undirected case is a subset of the set of SE in the directed case (up to cycles of cost 0). Therefore, upper bounds on the SPoA in games played on directed SP networks carry over to the undirected case. In addition, the examples that are used to show that these upper bounds are tight (or that the SPoA can be unbounded in general networks) hold when the edges are undirected as well. In the case of games with homogeneous capacities that are played on EP and SPP networks, we have shown that every SE is optimal, and this result carries over to the undirected case (by the above claim).

The upper bounds on the SPoS are derived from the fact that $SPoS \leq SPoA$, and we have shown that these bounds are asymptotically tight. The examples that are used to show that these upper bounds are tight still hold when the edges are undirected (the analysis of the game with unbounded SPoS appears in the appendix). This completes the extension of our results for symmetric games to the case of undirected graphs.

APPENDIX

A. OMITTED PROOFS

PROOF OF THEOREM 6.1. One direction has already been proven in Section 3: For every non-SPP network, we have shown that there is a symmetric CFCSC game played on it that does not admit a SE. We now show that every single source game played on an SPP network admits a SE.

Let Δ be a single source game played on an SPP network G . We denote the common source node by s and the sink node of agent j by t_j . By the definition of SPP networks, G is series composition of networks of parallel paths. Let G_1, \dots, G_k be networks of parallel paths such that $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$. For now, we assume that s is the source node of G_1 .

Consider the sink node of agent j . That node must be either an inner node or a sink node of one of the subnetworks G_1, \dots, G_k . Assume that t_j is an inner node or the sink of subnetwork G_i . Now let us look at a path from s to t_j . Since the graph is directed, any path from s to t_j does not use edges from subnetworks G_{i+1}, \dots, G_k . In addition, any path from s to t_j can be divided into i subsequent parts. The first $i - 1$ parts are paths from the source to the sink of each of the subnetworks G_1, \dots, G_{i-1} . The last part is a path from the source node of G_i to t_j . If t_j is an inner node of G_i , there is only one path from the source of G_i to t_j .

We have seen that the sink node of an agent determines whether edges from a subnetwork G_i can be part of her path. Without loss of generality, we assume that the agents are ordered in a way that agents who use more subnetworks of G_1, \dots, G_k come first. Formally, for every two agents j_1, j_2 , such that t_{j_1} is an inner node or the sink node of G_{i_1} and t_{j_2} is an inner node or the sink node of G_{i_2} , if $j_1 < j_2$, then (i) $i_1 \geq i_2$, and (ii) if t_{j_1} is an inner node and t_{j_2} is not an inner node, then $i_1 > i_2$.

Now we define a strategy profile \mathbf{r} and show that it is a SE. The definition of the profile uses Algorithm 1 in a similar way to the proof of Theorem 3.4. The profile \mathbf{r} is defined in two steps. First, if agent j 's sink node is an inner node of subnetwork G_i , every path from s to t_j uses the same edges in G_i , so we assign these edges to the agent. This is only a part of agent j 's path, and the rest of her path is determined in the next step. In the next step, we iterate through the subnetworks G_1, \dots, G_k . When subnetwork G_i is considered, we determine how many agents need to be assigned a path from the source node to the sink node of G_i and compute these paths using Algorithm 1. The algorithm is stated for edges, but as in the proof of Theorem 3.4, each path can be replaced by an equivalent edge that sums the costs of the edges in the path and has the minimal available capacity of these edges. Then, we assign each of the paths returned by the algorithm to the agents according to their order: The first agent j_1 who needs to be assigned a path in G_i gets the lowest-cost path returned by the algorithm. Then, the second agent $j_2 > j_1$ who needs a path in G_i is assigned the second lowest-cost path, and we continue the same way until all the paths returned by the algorithm are assigned. This completes the definition of \mathbf{r} .

We claim that the strategy profile \mathbf{r} is a SE. Assume by contradiction that there is a coalition C for which deviating to another profile is profitable, and let j be the minimal index of an agent in C . Agent j must reduce the cost that she pays for the edges in at least one of the subnetworks G_1, \dots, G_k .

For each G_i , agent j can reduce her cost in G_i by moving to another path or by sharing the cost of her path with another agent that deviates to that path. As in the proof of Theorem 3.4, agent j cannot use another path in a profitable deviation, as, if there were such a path, the algorithm would have assigned that path to j earlier. So after the deviation, agent j must share her path in G_i with more agents compared to the original profile \mathbf{r} . There are two cases:

- (1) Agent j 's path reaches the sink node of G_i . In that case, either agent j 's path in G_i has an edge e that is used by c_e agents, or the path is used by agents j, \dots, n . If the path contains an edge e that is used by c_e agents in \mathbf{r} , the path cannot be used by more than c_e agents after the deviation. If the path is already used by agents j, \dots, n , no other agents can join the path, since $j = \min C$. In both cases, the cost of agent j in G_i does not decrease.
- (2) Agent j 's sink node is an inner node of G_i . In that case, the only agents who can join agent j 's path in G_i are agents who reach the sink node of G_i . However, the agents are ordered such that the agents who reach the sink node of G_i are indexed lower than j . These agents do not participate in the deviation, and we conclude that agent j cannot reduce the cost of her path in G_i .

We have shown that the cost of agent j cannot decrease after the deviation, which contradicts the assumption that the deviation is profitable. Therefore, \mathbf{r} is a SE.

In the proof, we assumed that the source node s is the source node of G_1 . If s is the source node of another subnetwork G_i , the subnetworks G_1, \dots, G_{i-1} are never used by any agent and can be ignored. If s is an inner node of subnetwork G_i , all the possible paths of all agents use the same edges in G_i , and no agent can reduce the cost of her path in G_i . \square

PROOF OF THEOREM 6.2. We first prove that the networks described in the theorem admit a SE. The proof is very similar to that of Theorem 6.1. Assume that $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$, where G_l is a network of parallel paths and $G_1, \dots, G_{l-1}, G_{l+1}, \dots, G_k$ are networks of parallel edges. Let Δ be a multi-source game played on G .

As in the proof of Theorem 6.1, the source and sink nodes of each agent define which of the subnetworks G_1, \dots, G_k are used as part of her path. If the source or sink node of an agent is an inner node of G_l , there is only one path in G_l that reaches that node and the agent must use it.

In this proof, we assume that the agents are ordered in a way that agents whose source or sink node is an inner node of G_l come after all other agents. Formally, if j_1 is an agent that both her source and sink nodes are not inner nodes of G_l , and j_2 is an agent whose source or sink node is an inner node of G_l , then $j_1 < j_2$.

We now define a strategy profile \mathbf{r} and show that it is a SE. The profile \mathbf{r} is defined in two steps. First, if agent j 's source or sink node is an inner node of subnetwork G_l , every possible path for that agent uses the same edges in G_l , so we assign these edges to the agent. This is only a part of agent j 's path, and the rest of her path is determined in the next step. In the next step, we iterate through the subnetworks G_1, \dots, G_k (including G_l). When subnetwork G_i is considered, we determine how many agents need to be assigned a path from the source node to the sink node of G_i , and compute these paths using Algorithm 1. Then, we assign each of the paths returned by the algorithm to the agents according to their order: The first agent j_1 who needs to be assigned a path in G_i gets the lowest-cost path returned by the algorithm. Then, the second agent $j_2 > j_1$ who needs a path in G_i is assigned the second lowest-cost path, and so on. This completes the definition of \mathbf{r} .

We claim that the strategy profile \mathbf{r} is a SE. Assume by contradiction that there is a coalition C that benefits from deviating to another profile, and let j be the minimal index of an agent in C .

Agent j can reduce her cost in G_i by moving to another path or by sharing the cost of her path with another agent that deviates to that path. As in the proof of Theorem 3.4, agent j cannot use another path in a profitable deviation, as, if there were such a path, the algorithm would have assigned that path to j earlier. So after the deviation, agent j must share her path in G_i with more agents compared to the original profile \mathbf{r} . There are two cases:

- (1) Agent j 's path contains both the source and sink nodes of G_i . In that case, either agent j 's path in G_i has an edge e that is used by c_e agents, or the path is used by agents j, \dots, n . If the path contains an edge e that is used by c_e agents in \mathbf{r} , the path cannot be used by more than c_e agents after the deviation. If the path is already used by agents j, \dots, n , no other agents can join the path, since $j = \min C$. In both cases, the cost of agent j in G_i does not decrease.
- (2) $i = l$ and agent j 's source or sink node is an inner node of G_l . In that case, the only agents who can join agent j 's path in G_l are agents whose path contains both the source and sink nodes of G_l . However, the agents are ordered so that these agents are indexed lower than j . These agents do not participate in the deviation, and we conclude that agent j cannot reduce the cost of her path in G_l .

We have shown that the cost of agent j cannot decrease after the deviation, which contradicts the assumption that the deviation is profitable. Therefore, \mathbf{r} is a SE.

The converse direction relies on previous proofs. For every non-SPP network, we have already shown that there is a symmetric CFCSC game played on it that has no SE. Epstein et al. [2009] have provided an example of a game with multiple source and sink nodes that is played on the network depicted in Figure 11(a) and does not admit a SE. An equivalent game can be played on the network presented in Figure 11(b)

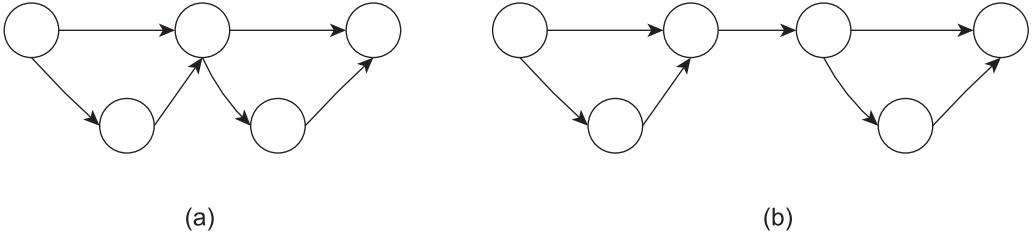


Fig. 11. SPP networks that do not admit a SE in games with multiple source and sink nodes.

by setting the cost of the additional edge to be 0. It is easy to show that one of the networks of Figure 11 is embedded in every SPP network $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$ that at least two of its subnetworks G_1, \dots, G_k are not networks of parallel edges. By the arguments used in the proof of Theorem 3.6, it follows that these networks do not admit a SE. \square

PROOF OF THEOREM 6.4. Let Δ be a game played on an SPP network G such that $G = G_1 \rightarrow G_2 \rightarrow \dots \rightarrow G_k$, and let \mathbf{s} be a SE in Δ . As in the proof of Theorem 4.1, it is sufficient to show that there is an optimal strategy profile \mathbf{s}^* , such that for every coalition of agents C , the profile $(\mathbf{s}_C, \mathbf{s}_{-C}^*)$ is feasible.

Let G_{OPT} be the subnetwork of G that is used by an optimal solution. We define a strategy profile \mathbf{s}^* that uses edges only from G_{OPT} . The source and sink nodes of each agent define which of the subnetworks G_1, \dots, G_k are used as part of her path. We iterate through G_1, \dots, G_k , and for every agent who uses edges from G_i , we define the path in G_i that this agent uses in \mathbf{s}^* .

First, for each of the parallel paths in G_i that are used in both \mathbf{s} and G_{OPT} , we assign each path to the same agents who use it in \mathbf{s} . This includes all the agents whose source or sink node is an inner node of G_i , who must use a specific path in G_i in every strategy profile. This does not violate the capacity constraints due to the feasibility of \mathbf{s} . Then, for the other agents who use G_i and still do not have a path, we assign them any available path in G_i that is part of G_{OPT} . Since there is a feasible strategy profile that uses only edges from G_{OPT} , there are enough available paths. This completes the definition of \mathbf{s}^* .

We now prove that for every coalition of agents C , the profile $\mathbf{s}_{comb} = (\mathbf{s}_C, \mathbf{s}_{-C}^*)$ is feasible. Let e be an edge in G_i . Denote by M the set of agents that use e in \mathbf{s} , and let M^* denote the set of agents that use e in \mathbf{s}^* . The set of agents that use e in \mathbf{s}_{comb} is $(M \cap C) \cup (M^* \cap (N \setminus C))$, where N is the set of all agents. If e is used only in \mathbf{s} , it cannot exceed its capacity in \mathbf{s}_{comb} , since $M \cap C \subseteq M$, and the profile \mathbf{s} is feasible. The same applies to edges that are used only in \mathbf{s}^* . It remains to consider edges that are used in both \mathbf{s} and \mathbf{s}^* . In the definition of \mathbf{s}^* , every agent who uses the simple path in G_i that contains e in \mathbf{s} is also assigned the same path in \mathbf{s}^* . It follows that $M \subseteq M^*$, and $(M \cap C) \cup (M^* \cap (N \setminus C)) \subseteq M^*$. Since \mathbf{s}^* is a feasible solution, we conclude that e does not exceed its capacity in \mathbf{s}_{comb} . Therefore, \mathbf{s}_{comb} is feasible. \square

PROOF OF THEOREM 5.2 FOR UNDIRECTED GRAPHS. The extension of the theorem to undirected graphs has two parts. First, we have to show that the SE described in the proof is still a SE. Recall that in this profile, one agent uses the path (s, e, c, d, b, t) and pays 0.6, and the other agent uses the path (s, a, b, t) and pays $R + 1.3$. Note that the lowest possible cost an agent can incur is still 0.6, which means that the agent who pays 0.6 will not participate in a deviation. The other agent cannot deviate to any path that excludes the edge of cost R . Hence, the profile is a SE.

Second, we have to show that when the edges are undirected, there are no new SE that exclude the edge (b, t) that costs R . Note that any path in a SE does not contain a cycle that costs more than 0. We analyze the strategy profiles that use edges in the opposite direction in order to look for new SE.

Using any of the edges that are adjacent to the source or the sink in the opposite direction results in a cycle of non-zero cost. Hence, this cannot occur in a SE.

Any path that uses the edge (c, e) (in that direction) either contains the edge (e, t) or the edge (e, s) . If an agent uses both (c, e) and (e, t) , the edge (s, e) cannot be used by the other agent, and deviating to the path (s, e, t) is profitable. If the path of an agent contains the subpath (c, e, s) , the path must contain a cycle, which is impossible in a SE.

Any feasible profile in which an agent uses one of the edges (b, a) , (c, a) , or (d, c) (in these directions) must contain a cycle that costs more than 0.

Finally, the only feasible profile without cycles in which the edge (b, d) is used (in that direction) contains the two paths (s, a, b, d, t) and (s, e, t) . The costs of the paths are 1.8 and 1, respectively, and the agents can benefit by deviating to another profile, such as (s, a, c, d, t) and (s, a, b, t) . We get that every SE must use the edge that costs R (and R can be arbitrarily high). \square

ACKNOWLEDGMENTS

Idan Nurick was involved in a preliminary version of this work. We thank Nick Gravin for helpful discussions.

REFERENCES

- Susanne Albers. 2009. On the value of coordination in network design. *SIAM J. Comput.* 38, 6 (March 2009), 2273–2302. DOI: <http://dx.doi.org/10.1137/070701376>
- Nir Andelman, Michal Feldman, and Yishay Mansour. 2009. Strong price of anarchy. *Games Econ. Behav.* 65, 2 (March 2009), 289–317. DOI: <http://dx.doi.org/10.1016/j.geb.2008.03.005>
- Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. 2008a. The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38, 4 (Nov. 2008), 1602–1623. DOI: <http://dx.doi.org/10.1137/070680096>
- Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. 2008b. Near-optimal network design with selfish agents. *Theor. Comput.* 4, 4 (Aug. 2008), 77–109. DOI: <http://dx.doi.org/10.4086/toc.2008.v004a004>
- Robert J. Aumann. 1959. Acceptable points in general cooperative n-person games. In *Contributions to the Theory of Games IV, Annals of Mathematical Study 40*, A. W. Tucker and R. D. Luce (Eds.). Princeton University Press, Princeton, NJ, 287–324.
- Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, and Gianpiero Monaco. 2013. Improved lower bounds on the price of stability of undirected network design games. *Theor. Comput. Syst.* 52, 4 (May 2013), 668–686. DOI: <http://dx.doi.org/10.1007/s00224-012-9411-6>
- Vittorio Bilò, Michele Flammini, and Luca Moscardelli. 2013. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS'13)*. IEEE Computer Society, Washington, DC, USA, 638–647. DOI: <http://dx.doi.org/10.1109/FOCS.2013.74>
- Steve Chien and Alistair Sinclair. 2009. Strong and pareto price of anarchy in congestion games. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part I (ICALP'09)*. Springer, Berlin, 279–291. DOI: http://dx.doi.org/10.1007/978-3-642-02927-1_24
- Amir Epstein, Michal Feldman, and Yishay Mansour. 2009. Strong equilibrium in cost sharing connection games. *Games Econ. Behav.* 67, 1 (Sept. 2009), 51–68. DOI: <http://dx.doi.org/10.1016/j.geb.2008.07.002>
- Thomas Erlebach and Matthew Radoja. 2015. Further results on capacitated network design games. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT'15)*. Springer, Berlin, 57–68. DOI: http://dx.doi.org/10.1007/978-3-662-48433-3_5
- Michal Feldman and Tom Ron. 2015. Capacitated network design games. *Theory of Computing Systems* 57, 3 (Oct. 2015), 576–597. DOI: <http://dx.doi.org/10.1007/s00224-014-9540-1>

- Amos Fiat, Haim Kaplan, Meital Levy, and Svetlana Olonetsky. 2007. Strong price of anarchy for machine load balancing. In *Proceedings of the 34th International Conference on Automata, Languages and Programming (ICALP'07)*. Springer, Berlin, 583–594. DOI: http://dx.doi.org/10.1007/978-3-540-73420-8_51
- Ron Holzman and Nissan Law-Yone. 1997. Strong equilibrium in congestion games. *Games Econ. Behav.* 21, 12 (Oct. 1997), 85–101. DOI: <http://dx.doi.org/10.1006/game.1997.0592>
- Ron Holzman and Nissan Law-Yone. 2003. Network structure and strong equilibrium in route selection games. *Math. Soc. Sci.* 46, 2 (Oct. 2003), 193–205. DOI: [http://dx.doi.org/10.1016/S0165-4896\(03\)00076-3](http://dx.doi.org/10.1016/S0165-4896(03)00076-3)
- Ron Holzman and Dov Monderer. 2015. Strong equilibrium in network congestion games: Increasing versus decreasing costs. *Int. J. Game Theor.* 44, 3 (Aug. 2015), 647–666. DOI: <http://dx.doi.org/10.1007/s00182-014-0448-4>
- Elias Koutsoupias and Christos Papadimitriou. 2009. Worst-case equilibria. *Comput. Sci. Rev.* 3, 2 (May 2009), 65–69. DOI: <http://dx.doi.org/10.1016/j.cosrev.2009.04.003>
- Jian Li. 2009. An $O(\log(n)/\log(\log(n)))$ upper bound on the price of stability for undirected Shapley network design games. *Inform. Process. Lett.* 109, 15 (July 2009), 876–878. DOI: <http://dx.doi.org/10.1016/j.ipl.2009.04.015>
- Igal Milchtaich. 2006. Network topology and the efficiency of equilibrium. *Games Econ. Behav.* 57, 2 (Nov. 2006), 321–346. DOI: <http://dx.doi.org/10.1016/j.geb.2005.09.005>
- Dov Monderer and Lloyd S. Shapley. 1996. Potential games. *Games Econom. Behav.* 14, 1 (May 1996), 124–143. DOI: <http://dx.doi.org/10.1006/game.1996.0044>
- Robert W. Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theor.* 2, 1 (Dec. 1973), 65–67. DOI: <http://dx.doi.org/10.1007/BF01737559>

Received August 2015; revised March 2016; accepted May 2016