Do Capacity Constraints Constrain Coalitions?*

Michal Feldman
Tel Aviv University
michal.feldman@cs.tau.ac.il

Ofir Geri
Tel Aviv University
ofirgeri@mail.tau.ac.il

Abstract

We study strong equilibria in symmetric capacitated cost-sharing 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 variants of cost-sharing games have been previously studied: (i) games where *coalitions* can form, and (ii) games where edges are associated with *capacities*; both variants are inspired by real-life scenarios. In this work we combine these variants 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 variants. Our contribution is two-fold. 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 are qualitatively different than those obtained in the analysis of each variant alone, and the combination of coalitions and capacities entails the introduction of more refined topology classes than previously studied.

1 Introduction

In recent years, a significant portion of AI research has departed from focusing on single agents to revolving around the study of multiagent systems. The construction of networks by autonomous agents is a typical example of this shift in focus. These situations can be modeled as *cost-sharing* connection games, which have been extensively studied (Anshelevich et al. 2003; 2004; Epstein, Feldman, and Mansour 2009; Albers 2009; Feldman and Ron 2012). 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 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 from which no agent can benefit by a unilateral deviation. In recent years, two interesting variants 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. SE in cost-sharing connection games have been studied in (Epstein, Feldman, and Mansour 2009). The second variant considers capacity constraints on the network edges (Feldman and Ron 2012). Here, each edge is associated with some capacity that limits the number of agents who can use it. Both variants 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 variants 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 variants.

Example 1 [Equilibrium existence]. In the plain setting (i.e., 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 the variants of coalitions and 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 2012). Consider next the combination of coalitions and capacities. Here, even the mere existence problem becomes non-trivial. Consider, for example, the game depicted

*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.

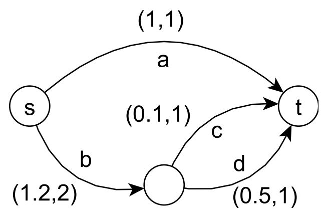

Figure 1: A game with no SE played on a simple EP network

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. One can easily verify that 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.

Example 2 [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). In the plain game, 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. 2003). 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 variant alone (see more details below).

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; i.e., 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 1 for a subset of our results). We provide tight bounds on the

		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

Table 1: A comparison between upper bounds for the PoA in different scenarios. Cap. and uncap. are shorthands for capacitated and uncapacitated. Our results are marked with an asterisk (*). All bounds are tight.

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 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.

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. Due to space constraints, these extensions appear in the full version (Feldman and Geri 2014).

1.2 Related Work

Cost-sharing connection games were introduced by Anshelevich et al. (2003), and have been widely studied since.

The fair cost-sharing mechanism has been studied in (Anshelevich et al. 2004). A key property of fair cost-sharing games is the fact that they admit an exact potential function and thus possess a pure strategy NE (Rosenthal 1973; Monderer and Shapley 1996). Settings with coalitions in congestion games have been studied by Holzman and Law-Yone (1997; 2003), and Epstein et al. (2009) studied the existence and quality of SE in cost sharing games. The characterization of network topologies that admit SE in congestion games with monotone cost functions has been later established by Holzman and Monderer (2014). The PoA measure has been introduced by Koutsoupias and Papadimitriou (1999) to study the quality of NE in games. The analogue of the PoA with respect to SE (called the strong PoA) was introduced and analyzed by Andelman et al. (2007). The SPoA of cost-sharing games was studied in (Epstein, Feldman, and Mansour 2009; Albers 2009), and was also studied with respect to various additional settings; see, e.g., (Fiat et al. 2007; Chien and Sinclair 2009). The consideration of capacities in cost-sharing games was first suggested by Feldman and Ron (2012).

2 Model and Preliminaries

2.1 Symmetric Capacitated Cost-Sharing Games

A symmetric capacitated cost-sharing (CCS) connection 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 $p_e \in \mathbf{R}^{\geq 0}$ and a capacity constraint $c_e \in \mathbf{N} \cup \{0\}$. Each agent wishes to construct an s-t path in G while maintaining minimal cost. The strategy space of agent i, denoted by Σ_i, is the set of all s-t paths in G. The joint strategy space is denoted by $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$.

Given a strategy profile $s = (s_1, \ldots, s_n) \in \Sigma$, the number of agents that use an edge e in the profile s is denoted by $x_e(s) = |\{i|e \in s_i\}|$. A profile s is said to be feasible if for every $e \in E$, $x_e(s) \leq c_e(s)$. A game is said to be feasible if it admits a feasible strategy profile. Throughout this paper we consider feasible games. For a given strategy profile s and a coalition C, the induced strategy profile on the agents of the coalition C is denoted by s_C, and the strategy profile of the rest of the agents is denoted by 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 i in a strategy profile s is

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

We use the utilitarian objective function, that is, the social cost of a strategy profile s is the sum of costs of all agents, $cost(s) = \sum_i p_i(s)$.

A strategy profile s is a Nash equilibrium (NE) if no agent can improve her cost by deviating to another strategy, i.e., for every i and every strategy $s'_i \in \Sigma_i$, it holds that $p_i(s) \leq p_i(s'_i, s_{-i})$ (where s_{-i} denotes the strategy profile of all agents except i in 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 s is a SE if for every coalition of agents C and every set of strategies $s'_C \in \Sigma_C$, there exists an agent $j \in C$ such that $p_j(s) \leq p_j(s'_C, s_{-C})$. 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 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 s^*, namely $PoA = \max_{s\in NE(\Delta)} \frac{cost(s)}{cost(s^*)}$. Similarly, the PoS is the ratio between the cost of the best-case NE and the cost of s^*, namely $PoS = \min_{s\in NE(\Delta)} \frac{cost(s)}{cost(s^*)}$. 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, $SPoA = \max_{s\in SE(\Delta)} \frac{cost(s)}{cost(s^*)}$ and $SPoS = \min_{s\in SE(\Delta)} \frac{cost(s)}{cost(s^*)}$. 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 $t \in V$. 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 $V' = (V \cup \{v\})\backslash\{v_1, v_2\}$ 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 $G = G_1 \to G_2$ 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 $G' = (V_1 \cup V_2, E_1 \cup E_2)$.

The following are classes of network topologies that will be of interest throughout the paper:

- A 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 so that $G = G_1 \to G_2$ or $G = G_1 \parallel G_2$.
- A network $G = (V, E)$ is an extension-parallel (EP) network if one of the following applies: (i) G consists of a single edge, (ii) There are two EP networks G_1, G_2 so

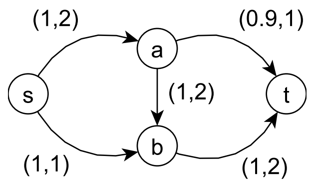

Figure 2: A game with no SE played on a Braess graph

that $G = G_1 \parallel G_2$, (iii) There is an EP network G_1 and an edge e so that $G = G_1 \to e$ or $G = e \to G_1$.

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: (i) Subdivision (replacing an edge (u, v) by a new node w and two edges (u, w) and (w, v)), (ii) Addition of a new edge connecting two existing nodes (including nodes that were added using subdivision or extension), (iii) 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, Feldman, and Mansour 2009). In the capacitated version, it has been shown by Feldman and Ron (2012) 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 in a simple EP graph.

An additional example of a game that does not admit a SE is depicted in Figure 2, on an underlying network known as a Braess graph. It is not difficult to verify that a game with two agents played on this graph does not admit any SE either.

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. We first introduce a new family of network topologies.

Definition 3.1. A symmetric network G is said to be *Series of Parallel Paths* (SPP) if there exist networks $G_1, \ldots, G_k$, each constructed by a parallel composition of simple paths, such that $G = G_1 \to \ldots \to G_k$.

Note that every SPP network is an SP network, but may or may not be an EP network.

Definition 3.2. A symmetric network G is said to admit a SE if *every* symmetric CCS 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* CCS 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. For every edge e that has not been assigned agents yet, let its *fractional cost* be $\frac{p_e}{\min\{c_e, n - \sum_{i=1}^{j-1} n_i\}}$.
3. Assign the edge e with the minimal fractional cost to $n_j = \min\{c_e, n - \sum_{i=1}^{j-1} n_i\}$ agents.
4. If there are agents that have not been assigned an edge, increment j and go back to step 2.

Theorem 3.3. *A 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.

Let $G_1, \ldots, G_k$ be parallel-edge networks, and let G be a series composition of these networks. Algorithm 1 computes a strategy profile for games played on parallel-edge networks. The algorithm first computes the lowest-cost path in the network, and assigns it to n_1 agents (n_1 is calculated by the algorithm). In the j^{th} iteration, the lowest-cost path that is still available is determined, and assigned to n_j agents (n_j is calculated by the algorithm). Since the game is symmetric, Algorithm 1 is only concerned with the number of the agents (denoted n_j) and not with their identity. After the number of agents that use each edge is decided, the actual agents can be assigned strategies. The algorithm terminates within at most n steps. In this proof, we show that the returned profile is a SE. For every $i = 1, \ldots, k$, we compute a profile s_i for network G_i using Algorithm 1. Since the game is symmetric we may reorder the agents in an increasing order of their costs, namely, for every profile s_i, we order the agents so that for every $j < j'$, $p_j(s_i) \leq p_{j'}(s_i)$.

Let s be the strategy profile in which the strategy of agent j is the concatenation of his strategies in $s_1, \ldots, s_k$. Since we sorted each of the profiles s_i, the first agent in s uses the lowest-cost path in each of $G_1, \ldots, G_k$, the next agent uses the second lowest-cost path, and so on. We wish to prove that s is a SE. Assume that there is a coalition C for which there is a profitable deviation, yielding a strategy profile s'. Let j be the minimal index of an agent in C, i.e., j is the agent that uses the j^{th} lowest-cost path in each of $s_1, \ldots, s_k$.

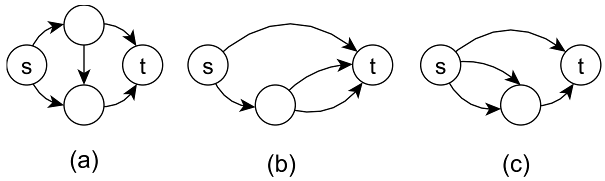

Figure 3: Minimal non-SPP networks

In order for s' to be a profitable deviation, j must reduce his cost in one of $G_1, \ldots, G_k$. Assume, without loss of generality, that by deviating to s', agent j reduces his cost in G_1. Denote the edge used by j in s_1 by e. There exists an edge e' so that

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

The edge e' cannot be one of the edges entirely used in the profile s_1. If e' is an edge that is not used in s_1, it holds that $x_{e'}(s') \leq \min\{c_{e'}, n - j + 1\}$. If e' is partially used in s_1, either $e' = e$ (in that case, $x_{e'}(s') \leq x_e(s)$, and j cannot reduce his cost) or $x_{e'}(s') \leq \min\{c_{e'}, n - j + 1\}$. Therefore,

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

and the algorithm should have assigned e' to agent j (or a previous agent) instead of e. We get a contradiction, thus, s is a SE. □

We now prove the converse direction, namely that for every non-SPP network, there is a CCS 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. The proof appears in the full version (Feldman and Geri 2014).

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

Using Lemma 3.5 we can now prove the following. The proof is similar to that of a lemma proven in (Feldman and Ron 2012), and is specified here for completeness.

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

Proof (Sketch). 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. For every edge in the corresponding network, assign costs and capacities as in Figures 1 and 2. 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 $p_{e_2} = 0$. If the source or the sink are extended using a new edge e, we set $c_e = 2$ and $p_e = 0$. If an edge e connecting two existing nodes is added, we set its capacity to be $c_e = 0$. This construction ensures that each strategy profile in the game played on G emulates a strategy profile in one of the examples presented earlier in this section, where a SE does not exist. □

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 strong price of anarchy (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 CCS game played on an EP network, it holds that $SPoA \leq H_n$ (if a SE exists).*

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 s that is a SE (by definition). In particular, if some coalition C deviates to its corresponding profile in the socially optimal profile s^*, then one of the agents in C weakly prefers the initial profile s to the new profile (s^*_C, s_{-C}). The desired bound is then derived by the obtained inequalities for coalitions of sizes $n, \ldots, 1$. The only barrier to applying the exact same technique to capacitated games is the fact that the deviation into profile (s^*_C, 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. Note that there may be more than a single socially optimal profile, but we have the freedom to choose one of them to be s^*.

Lemma 4.2. *Let G be an EP network and s be a SE in a symmetric CCS game played on G. There exists a feasible strategy profile s^* so that the cost of s^* is minimal, and for every coalition C, the profile (s^*_C, s_{-C}) is feasible.*

The following lemma will be used in the proof.

Lemma 4.3. *(Feldman and Ron 2012) Let G be an SP network. Let s be a feasible profile of k agents in a game played on G, and let s' be a feasible profile of r agents such that $r < k$. There exists an $s - t$ path in G that is feasible in s' and uses only edges that are used in the profile 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 s^* played on G_{OPT}, and then prove that for every coalition C, the strategy profile (s^*_C, s_{-C}) is feasible. Since G is EP, G_{OPT} is also EP. We define 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 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 s, which is a SE.

It is important to note that the algorithm might possibly define a strategy profile only for a subset of the agents. In step 2, $N_1 \cap N_2 = \phi$, 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 the agents in N will be assigned a strategy. 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 Lemma 4.3 to

Algorithm 2 Choosing the optimal profile s^*

Input: G_{OPT} is a graph, N is a set of agents, and s is a strategy profile.
ChooseOptimalProfile(G_{OPT}, N, s):

1. If $G_{OPT} = e$, where e is a single edge, return a strategy profile in which all of the agents in N use the edge e.
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) $s_1 \leftarrow ChooseOptimalProfile(G_1, N_1, s)$
   (c) $s_2 \leftarrow ChooseOptimalProfile(G_2, N_2, s)$
   (d) Return the union of the profiles: (s_1, s_2).
3. If $G_{OPT} = G_1 \rightarrow e$ or $G_{OPT} = e \rightarrow G_1$, where e is an edge:
   (a) $s_1 \leftarrow ChooseOptimalProfile(G_1, N, s)$
   (b) Each agent that has a strategy in s_1 will also use the edge e.
   (c) For any other agent that uses e in 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 s.

the network G_{OPT}, where s is any strategy profile that uses only edges from G_{OPT}, and s' is the partial profile that was defined above. This provides a full definition of the profile s^*.

We claim that the profile 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 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 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 wish to prove that the strategy profile $s_{comb} = (s_C, s^*_{-C})$ is feasible. Let e be an edge. Let M denote the set of agents that use e in s, and let M^* denote the set of agents that use e in s^*. The set of agents that use e in s_{comb} is $(M \cap C) \cup (M^* \cap (N \backslash C))$. If e is used only in s, it cannot exceed its capacity in s_{comb}, since $M \cap C \subseteq M$, and the profile s is feasible. The same applies to edges that are used only in s^*. It remains to consider edges that are used in both s and 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 s and s^*, there are two cases:

Case 1: If e is assigned agents in step 1 of the algorithm, then the set N in this step contains all the agents that use e in s, namely, $N = M$, thus, $M \subseteq M^*$ $(e \in s_i \Rightarrow e \in s^*_i)$.

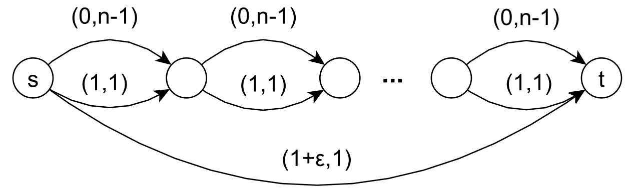

Figure 4: A game played on an SP network with $SPoA = n$

The set of agents that use e in s_{comb} is $(M \cap C) \cup (M^* \cap (N \backslash C)) \subseteq M^*$. Since s^* is feasible, then e does not exceed its capacity in s_{comb}.

Case 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 cases. In step 3(c), if each agent that uses e in s was assigned a path in G_1 in s^*, we get that $M \subseteq M^*$. The set of agents that use e in s_{comb} is $(M \cap C) \cup (M^* \cap (N \backslash C)) \subseteq M^*$. Since s^* is feasible, we get that e does not exceed its capacity in s_{comb}. The second case is that there is an agent that uses e in s but does not use e and G_1 in s^*. Due to the definition of 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 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 s_{comb} is $(M \cap C) \cup (M^* \cap (N \backslash C)) \subseteq M$. Since s is feasible, then e does not exceed its capacity in s_{comb}.

We conclude that no edge exceeds its capacity in s_{comb}; the assertion of Theorem 4.1 follows. ☐

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

Theorem 4.4. *For every symmetric CCS game played on a network that is a series composition of EP networks, it holds that $SPoA \leq H_n$ (if a SE exists), and this bound is tight.*

The proof appears in the full version (Feldman and Geri 2014). 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 is established in (Feldman and Ron 2012) that $PoA \leq n$, which directly implies that $SPoA \leq n$. We provide an example showing that this bound is tight.

Theorem 4.5. *For every $\epsilon > 0$, there exists an SP network G and a CCS 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 4. In this game, $SPoA \geq \frac{n}{1+\epsilon}$. The analysis appears in the full version (Feldman and Geri 2014). ☐

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

Theorem 4.6. *For every real number R, there exists a CCS game with two agents in which the SPoA is greater than R.*

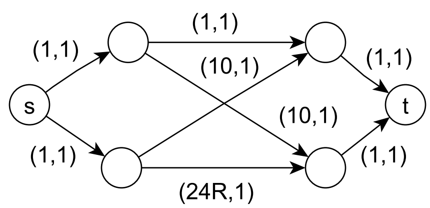

Figure 5: A game with unbounded SPoA

Proof. Consider the game played by two agents on the graph presented in Figure 5. In this game, $SPoA > R$. The analysis appears in the full version (Feldman and Geri 2014). □

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 $SE(\Delta) \subseteq NE(\Delta)$), 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. Due to space constraints, the proofs appear in the full version (Feldman and Geri 2014).

For EP and SPP networks, the bound of H_n is tight with respect to the best SE due to an example provided in (Feldman and Ron 2012), in which $PoS = \frac{H_n}{1+\epsilon}$ for every $\epsilon > 0$.

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 CCS game played on an SP network, in which the SPoS is at least* $\Omega(n)$.

Next, we show that as the SPoA, the SPoS can also be unbounded.

Theorem 5.2. *There exists a symmetric CCS game in which the SPoS can be arbitrarily high.*

Acknowledgments

Idan Nurick was involved in a preliminary version of this work.

References

Albers, S. 2009. On the value of coordination in network design. *SIAM J. Comput.* 38(6):2273–2302.

Andelman, N.; Feldman, M.; and Mansour, Y. 2007. Strong price of anarchy. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, 189–198.

Anshelevich, E.; Dasgupta, A.; Tardos, É.; and Wexler, T. 2003. Near-optimal network design with selfish agents. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, 511–520.

Anshelevich, E.; Dasgupta, A.; Kleinberg, J.; Tardos, É.; Wexler, T.; and Roughgarden, T. 2004. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '04, 295–304.

Aumann, R. J. 1959. Acceptable points in general cooperative n-person games. In Tucker, A. W., and Luce, R. D., eds., *Contributions to the Theory of Games IV, Annals of Mathematical Study 40*. Princeton University Press. 287–324.

Chien, S., and Sinclair, A. 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, 279–291.

Epstein, A.; Feldman, M.; and Mansour, Y. 2009. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior* 67(1):51–68.

Feldman, M., and Geri, O. 2014. Do capacity constraints constrain coalitions? arXiv.

Feldman, M., and Ron, T. 2012. Capacitated network design games. In *Proceedings of the 5th International Symposium on Algorithmic Game Theory*, SAGT '12, 132–143.

Fiat, A.; Kaplan, H.; Levy, M.; and Olonetsky, S. 2007. Strong price of anarchy for machine load balancing. In *Proceedings of the 34th International Conference on Automata, Languages and Programming*, ICALP '07, 583–594.

Holzman, R., and Law-Yone, N. 1997. Strong equilibrium in congestion games. *Games and Economic Behavior* 21(12):85 – 101.

Holzman, R., and Law-Yone, N. 2003. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences* 46(2):193 – 205.

Holzman, R., and Monderer, D. 2014. Strong equilibrium in network congestion games: increasing versus decreasing costs. *International Journal of Game Theory*.

Koutsoupias, E., and Papadimitriou, C. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, STACS '99, 404–413.

Monderer, D., and Shapley, L. S. 1996. Potential games. *Games and Economic Behavior* 14(1):124 – 143.

Rosenthal, R. W. 1973. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1):65–67.