Theory Comput Syst (2015) 57:576–597
DOI 10.1007/s00224-014-9540-1

Capacitated Network Design Games

Michal Feldman · Tom Ron

Published online: 8 April 2014

Abstract We study a *capacitated* symmetric network design game, where each of n agents wishes to construct a path from a network's source to its sink, and the cost of each edge is shared equally among the agents using it. The uncapacitated version of this problem has been introduced by Anshelevich et al. (2003) and has been extensively studied. We find that the consideration of edge capacities entails a significant effect on the quality of the obtained Nash equilibria (NE), under both the utilitarian and the egalitarian objective functions, as well as on the convergence rate to an equilibrium. The following results are established. First, we provide bounds for the price of anarchy (PoA) and the price of stability (PoS) measures with respect to the utilitarian (i.e., sum of costs) and egalitarian (i.e., maximum cost) objective functions. Our main result here is that unlike the uncapacitated version, the network topology is a crucial factor in the quality of NE. Specifically, a network topology has a bounded PoA if and only if it is *series-parallel* (SP), i.e., a network that is built inductively by series compositions and parallel compositions of SP networks. Second, we show that the convergence rate of best-response dynamics (BRD) may take $\Omega(n^{1.5})$ steps. This is in contrast to the uncapacitated version, where convergence is guaranteed within at most n iterations.

This work was partially supported by Israel Science Foundation (grant 1219/09), Google Inter-university Center for Electronic Markets and Auctions, and the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 274919.

M. Feldman (✉)
Blavatnic School of Computer Science, Tel Aviv University, Tel Aviv, Israel
e-mail: mfeldman@tau.ac.il

T. Ron
School of Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel
e-mail: rontom@gmail.com

Keywords Network design games · Game theory · Cost sharing · Capacity · Price of anarchy · Price of stability

1 Introduction

The construction of large networks by strategic agents has been widely studied from a game-theoretic perspective in the past decade [4, 11, 12, 30]. For a motivating example, consider the construction and maintenance of large computer networks by independent economic agents with different, and often competing, self-interests. The game-theoretic perspective offers tools and insights that are fundamental to the understanding and analysis of these settings.

In a symmetric network design game, a network is given, where each edge is associated with some cost; and a set of n agents wish to buy some path from the network's source (s) to its sink (t). Every agent chooses an s-t path, and the cost of every edge is divided equally among the agents who use it. This is often called a *fair cost-sharing* method. The game-theoretic twist is the assumption that each agent chooses its path strategically, so as to minimize its cost. It is well known that the Nash equilibria of this game need not be efficient, where efficiency is usually defined with respect to either the sum of the agents' costs (referred to as the *utilitarian* or *sum-cost* objective) or to the maximum cost of any agent (referred to as the *egalitarian* or *max-cost* objective).

The efficiency loss is commonly quantified using the price of anarchy (PoA) [21, 27] and price of stability (PoS) [4] measures; the former refers to the ratio between the cost of the worst Nash equilibrium and the social optimum, whereas the latter refers to the ratio between the cost of the best Nash equilibrium and the social optimum. The network design game described above is fairly easy to analyze. The PoA is known to be tightly bounded by n with respect to the utilitarian objective function[1] [4]. It is not too difficult to see that the same bound holds with respect to the egalitarian objective. In addition, the PoA is independent of the network topology, as the worst case is obtained for two parallel links. The PoS, in contrast, is always equal to 1 (with respect to both objective functions), since in a symmetric network, the profile in which all agents share the shortest path from s to t is a Nash equilibrium. Finally, best-response dynamics (i.e., dynamics in which agents sequentially apply their best-response moves) exhibits a simple structure, where convergence to a NE is guaranteed within at most n steps.

Interestingly, the majority of those results rely on the assumption that the network edges are *uncapacitated*; i.e., it is assumed that edges may hold any number of agents. While this assumption has been employed by most of the studies on strategic network formation games, we claim that in real-life applications, network links have a limit on the number of agents they can serve. To reflect this observation, we introduce *capacitated* network design games, in which every edge, in addition to its cost, is also associated with a *capacity* that specifies the number of agents it can support.

[1] While [4] consider an underlying directed graph, this bound carries over to the undirected case.

We study the quality of NE in these games (using both PoA and PoS measures) and the convergence rate of best-response dynamics. We are particularly interested in the effect of the *topology* of the underlying network on the obtained results.

In cases where edges are associated with capacities, a *feasibility* problem arises (i.e., whether there exists a solution that accommodates all the agents). However, as already hinted at by [4], if a feasible solution exists, the arguments used in the uncapacitated version can be applied to show that a pure NE exists and, moreover, every best-response dynamics converges to a pure NE. This observation motivates our study.

1.1 Our Contribution

For the PoA, the lower bound of n trivially carries over to the capacitated version; thus, one cannot expect a bound better than n. The upper bound of n, however, does not carry over. In particular, we demonstrate that the PoA can be arbitrarily high. As it turns out, the network topology plays a major role in the obtained PoA. A symmetric network topology G is said to be *PoA-bounded* if for every symmetric network design game that is played on G, the PoA is bounded by n, independently of the edge costs and capacities. Our main result here is a full characterization of PoA-bounded network topologies. Specifically, we show that a symmetric network topology is PoA-bounded if and only if it is a *series-parallel* (SP) network, i.e., a network that is built inductively by series compositions and parallel compositions of SP networks. This result holds with respect to both the sum-cost and max-cost objectives. Moreover, for parallel-link networks, we show that the PoA (with respect to both the sum-cost and max-cost objectives) is essentially bounded by the maximum edge capacity in the network, and that this bound is tight.

This separation between the graph topology and the assignment of edge costs and capacities reflects a separation between the underlying infrastructure and the edge characteristics. While the infrastructure is often stable over time, the edge characteristics may be modified over short time periods. A PoA-bounded topology ensures that, no matter how edge characteristics evolve, the cost of a NE will never exceed n. Such topologies are desired by network designers, who wish to guarantee efficiency in their network despite the fact that they do not control the actions of the individual users. Notably, within the class of SP networks, the worst case is obtained already for parallel links.

In contrast to the PoA, the PoS with respect to the sum-cost objective is not affected by the network topology. In particular, we provide a lower bound of $H(n)$ $\left(\text{i.e., } \sum_{i=1}^{n} \frac{1}{n}\right)$ for the PoS on parallel-link graphs, and show that for every symmetric network the PoS is upper bounded by $H(n)$ (this is a direct consequence of [4]). As for the max-cost objective function, for SP graphs the upper bound of n that is established for the PoA trivially carries over to the PoS, and a matching lower bound is established. For general graphs, we establish an upper bound of $n \log n$. Closing the gap between n and $\log n$ for the PoS in general graphs remains an open problem.

Finally, we study the efficiency loss with respect to the sum-cost objective under the special case of *homogeneous capacities*, i.e., where all the edges have the same capacity. Surprisingly, homogeneity of edge capacities has a very different effect on the PoS that on the PoA. Specifically, homogeneity ensures the existence of an optimal Nash equilibrium (i.e., PoS $= 1$), but the price of anarchy may still be arbitrarily high.

Most of our results regarding the PoA and PoS bounds are summarized in Table 1, where they are also contrasted with the corresponding results in the uncapacitated version (specified in parentheses). These results suggest that the departure from the classic assumption of uncapacitated edges results in significant differences in the quality of equilibria.

Additionally, the consideration of capacities introduces additional complexity that reveals itself through a slower conversion rate. While BRD in the uncapacitated version is guaranteed to converge within at most n iterations, we establish a lower bound of $\Omega(n^{3/2})$ for convergence in capacitated games. Moreover, this lower bound is obtained already in parallel link graphs.

1.2 Related Work

Various models of network design and network formation games have been extensively studied in the past decade from a game-theoretic perspective [4–6, 8, 10, 22], with a great emphasis on the PoA and PoS measures. The PoA in network design games has been also studied with respect to the *strong equilibrium* solution concept by Epstein et al. [11], Andelman et al. [3] and Albers [2].

Korillis et al. [20, 33] previously considered the influence of capacities on network games. In [20] Korillis et al. examined the problem of optimal capacity allocation under noncooperative routing. The network designer's aim is to allocate link capacities so that the resulting Nash equilibria are efficient according to some aggregative performance criterion. Moreover, it is shown that contrary to the common intuition, adding link capacity may lead to degradation of user performance. In [33] Korillis et al. proposed methods for efficiently adding resources to a non-cooperative network of general topology that guarantee an improvement in performance, thus establishing a methodology for efficiently coping with the Braess paradox [7] in non-cooperative networks.

Table 1 Summary of our results. All the results, except for the PoS w.r.t. max-cost for general networks, are tight. Our results are contrasted with the well-known bounds for uncapacitated games, which are specified in parentheses

		Parallel links	SP	General
sum-cost (sc)	PoA	n (n)	n (n)	unbounded (n)
	PoS	$\log n(1)$	$\log n(1)$	$\log n(1)$
max-cost (mc)	PoA	n (n)	n (n)	unbounded (n)
	PoS	n (1)	n (1)	$n \log n(1)$

The effect of edge capacities in network congestion games was also studied by Correa et al. [9], who extended the efficiency bounds in network routing games (with non-atomic traffic) [29, 30] to capacitated networks. They showed that the consideration of edge capacities leads to the existence of multiple equilibria, and that while the worst equilibrium is not necessarily nearly optimal, the best equilibrium is as efficient as the (unique) equilibrium in the equivalent network game without capacities.

Motivated by the fact that pure Nash equilibria in congestion games may be inefficient even in parallel-link networks in congestion games, Von Falkenhausen and Harks [32] studied the design of cost-sharing methods that improve the PoA. They focused on job-scheduling problems on parallel machines with non-decreasing cost functions, and characterized a class of instances that admit a bounded PoA.

The role that network topology plays in game-theoretic settings has been studied in various models. In the model of network routing, it has been shown by Roughgarden and Tardos [30] that the PoA is independent of the network topology. In contrast, the network topology seems to matter a lot in other settings. Some prominent examples include the following. Milchtaich [25] showed that the *Pareto efficiency* of equilibria in network routing games (with a continuum of agents) strongly depends on the network topology. In addition, topological characterizations for symmetric network games have been also provided for other equilibrium properties, including (Nash and strong) equilibrium existence (see Milchtaich [24], Epstein et al. [11, 12], and Holzman and Law-Yone [18, 19]), and equilibrium uniqueness (see Milchtaich [23]). Milchtaich [24] identified the topological conditions guaranteeing the existence of at least one pure Nash equilibrium in every network congestion game with player-specific costs or weights.

Holzman and Law-Yone [18, 19] characterized the network topologies where strong equilibrium (where no coalition can improve the cost of each of its members) always exists in monotone decreasing congestion games. They showed that a strong equilibrium always exists in the case where all strategies are singletons. In addition, they showed that in single-commodity networks a strong equilibrium exists if and only if the underlying graph is *extension parallel*.

Epstein et al. [11] studied graph topologies that guarantee the existence of a strong equilibrium in *fair* cost-sharing connection games and *general* cost-sharing connection games. It has been shown that (1) single-commodity networks always admit a strong equilibrium (in both fair and general games), (2) a strong equilibrium always exists in single-source multiple-sinks SP networks (in both fair and general games), and (3) a multi-source and sink network with an underlying *extension-parallel* graph always admits a strong equilibrium in fair connection games. They also examined the quality of the obtained strong equilibrium and established bounds on the PoA with respect to a strong equilibrium.

Epstein et al. [12] examined the topological structure of networks that guarantee that any Nash equilibrium achieves the social optimum. Their main contribution is in showing that for symmetric single-commodity congestion games, the topologies that guarantee this property are extension-parallel graphs, and for the family of *bottleneck* routing games, the topologies that guarantee this property are series-parallel graphs. In multi-commodity games it is shown that the efficient topologies are very limited (and include either trees or trees with parallel edges).

Best-response dynamics and its convergence rate has been the subject of intensive research recently. Since every congestion game is a potential game [26, 28], best-response dynamics always converge to a pure NE. However, best-response dynamics may in general be exponentially long, as established in a series of papers [1, 14, 31]. This observation has led to a large amount of work that identified special classes of congestion games, where best-response dynamics converge to a Nash equilibrium in polynomial time (or even linear time). This agenda was studied in settings with negative congestion effects (e.g., [13, 16]), settings with positive congestion effects (e.g., [4]), and settings with combined congestion effects (e.g., [15]). Finally, the robustness of best-response convergence to altruistic agents has been studied in [17], where it has been shown that BRD may cycle as a result of altruism.

2 Model and Preliminaries

2.1 Capacitated Symmetric Cost-Sharing Games

A capacitated symmetric cost-sharing connection (CCS) game (also known as single commodity) is a tuple

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

where n is the number of agents and $G = (V, E)$ is an undirected graph, with $s, t \in V$ as its *source* and *sink* nodes, respectively. Every edge $e \in E$ is associated with a cost $p_e \in R^{\geq 0}$ and a capacity $c_e \in N$, where an edge capacity specifies the maximum number of agents that can use it. The set of agents $\{1, \dots, n\}$ is also denoted by $[n]$. Every agent i wishes to construct an s-t path in G. The action space of an agent i, denoted by Σ_i, is the set of $s - t$ paths in G, and an action of an agent i (which is simply an $s - t$ path in G) is denoted by $S_i \in \Sigma_i$. Since this is a symmetric game, all agents have the same action space. The joint action space (indicating an $s - t$ path for each agent) is denoted by Σ.

We consider the *fair* cost-sharing game, where an edge's cost is shared equally by all the agents that use it in their path. Given an action profile $S = (S_1, \dots, S_n)$, we denote by $x_e(S)$ the number of agents that use edge e in their path; i.e., $x_e(S) = |\{i : e \in S_i\}|$. A profile S is said to be *feasible* if for every $e \in E$, $x_e(S) \leq c_e$. The cost of agent i in a profile S is defined as

$$p_i(S) = \begin{cases} \sum_{e \in S_i} \frac{p_e}{x_e(S)} & \text{, if } x_e(S) \leq c_e \text{ for every } e \in S_i \\ \infty & \text{, otherwise} \end{cases} \tag{1}$$

A profile S is said to be a *Nash equilibrium* if no agent can improve its cost by a unilateral deviation; i.e., for every i, $S'_i \in \Sigma_i$, $S_{-i} \in \Sigma_{-i}$, it holds that $p_i(S) \leq p_i(S'_i, S_{-i})$, where S_{-i} denotes the joint action of all agents except i.

Given a game Δ, let $\tau(\Delta)$ denote the set of all feasible profiles in Δ. A CCS game Δ is said to be feasible if it admits a feasible profile; i.e., $\tau(\Delta) \neq \emptyset$.

We consider two social cost functions. The *sum-cost* of a profile S is the total cost of the agents in S (and also equals the total cost of the purchased edges in S), and is given by

$$sc_{\Delta}(S) = \sum_{i} p_i(S)$$

The *max-cost* of a profile S is the maximum cost of any agent in S, and is given by

$$mc_{\Delta}(S) = \max_{i \in [n]} p_i(S)$$

We denote by $OPT_{sc}(\Delta)$ and $OPT_{mc}(\Delta)$ the optimal profiles with respect to the sum-cost and max-cost objectives, respectively. When clear in the context, we omit Δ, and also abuse notation and use OPT_{sc} and OPT_{mc} to denote the cost of the respective optimal solutions.

In the figures of the paper, every edge is associated with a tuple (c_e, p_e), denoting its capacity and cost, respectively.

2.2 Nash Equilibrium Existence

An uncapacitated fair cost-sharing game is known to be a *potential game* [4, 28]. Every potential game admits a pure NE [26, 28]. Moreover, BRD (where agents sequentially apply their best-response moves) always converge to a pure NE. Capacitated versions are not guaranteed to admit a feasible solution; however, if a feasible solution exists, then so does a pure NE.

Observation 1 [4] *Let Δ be a CCS game s.t. $\tau(\Delta) \neq \emptyset$. Then, Δ admits a pure NE and every best- response dynamics converges to a NE.*

This proof relies on the existence of a potential function, $\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{x_e(S)}$, that emulates the cost of an agent when deviating from one feasible solution to another.

2.3 Efficiency Loss

To quantify the efficiency loss due to strategic behavior, we use the PoA and PoS measures. The PoA is the ratio of the worst Nash equilibrium and the social optimum, and is given by $PoA_{sc}(\Delta) = \frac{\max_{S \in NE(\Delta)} sc_{\Delta}(S)}{OPT_{sc}(\Delta)}$ and $PoA_{mc}(\Delta) = \frac{\max_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)}$ with respect to the sum-cost and max-cost objectives, where $NE(\Delta)$ denotes the set of NE of Δ, and it is assumed that $NE(\Delta) \neq \emptyset$. Similarly, the PoS is given by $PoS_{sc}(\Delta) = \frac{\min_{S \in NE(\Delta)} sc_{\Delta}(S)}{OPT_{sc}(\Delta)}$ and $PoS_{mc}(\Delta) = \frac{\min_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)}$ with respect to the two objective functions, respectively.

2.4 Graph-Theoretic Preliminaries

In this section we provide some preliminaries regarding network topologies. A *symmetric* network is an undirected graph G along with two distinguished nodes, a source s and a sink t. When clear in the context, we refer to G as the symmetric network. A CCS game is symmetric (also called single-commodity) if its underlying network is symmetric with source s and sink t, and nodes s and t are the respective source and sink of all the agents. A symmetric network G is *embedded* in a symmetric network G' if G' is isomorphic to G or to a network derived from G by applying the following operations any number of times in any order: (i) *Subdivision* of an edge (i.e., its replacement by a path of edges), (ii) *Addition* of a new edge joining two existing nodes, (iii) *Extension* of the source or the sink (i.e., addition of a new edge joining s or t with a new node, which becomes the new source or sink, respectively).

Next, we define the following operations on symmetric networks:

- Identification: The *identification* operation is the collapse of two nodes into one. More formally, given a graph $G = (V, E)$ we define the *identification* of nodes $v_1 \in V$ and $v_2 \in V$ forming a new vertex $v \in V$ as creating a new graph $G' = (V', E')$ where $V' = V \setminus \{v_1, v_2\} \cup \{v\}$ and E' includes the edges of E where the edges of v_1 and v_2 are now connected to v.
- Parallel composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, with sources $s_1 \in V_1$ and $s_2 \in V_2$ and sinks $t_1 \in V_1$ and $t_2 \in V_2$, respectively, we define a new symmetric network $G = G_1||G_2$ as follows. Let $G' = (V_1 \cup V_2, E_1 \cup E_2)$ be the union network. To generate $G = G_1||G_2$ we identify the sources s_1 and s_2, forming a new source node s, and identify the the sinks t_1 and t_2, forming a new sink t.
- Series composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, with sources $s_1 \in V_1$ and $s_2 \in V_2$ and sinks $t_1 \in V_1$ and $t_2 \in V_2$, respectively, we define a new symmetric network $G = G_1 \to G_2$ as follows. Let $G' = (V_1 \cup V_2, E_1 \cup E_2)$ be the union network. To generate $G = G_1 \to G_2$ from G' we identify the vertices t_1 and s_2, forming a new vertex u. The network G has a source $s = s_1$ and a sink $t = t_2$.

A **series-parallel (SP)** network is a symmetric network that is constructed inductively from two SP networks by either a series composition or a parallel composition, where a single edge serves as the base of the induction. That is, a symmetric network consisting of a single edge is a SP network. In addition, given two SP networks, G_1 and G_2, the networks $G = G_1||G_2$ and $G = G_1 \to G_2$ are SP networks.

3 The Sum-Cost Objective Function

3.1 Price of Anarchy (PoA)

Throughout this section, we write PoA to denote PoA_{sc} for simplicity. In uncapacitated cost-sharing games, the PoA is n (tightly). This is, however, not the case in capacitated games, as demonstrated by the following proposition.

Proposition 2 *The price of anarchy with respect to the sum-cost function in CCS games can be arbitrarily high.*

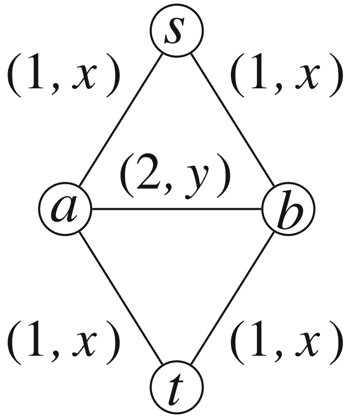

Fig. 1 The price of anarchy can be arbitrarily high

Proof Consider a CCS game with two agents and an underlying graph as depicted in Fig. 1, and suppose that y is arbitrarily larger than x. The optimal profile is where one agent uses the path s-a-t and the other uses the path s-b-t, resulting in a total cost of $4x$. However, there is a NE in which one agent uses the path s-a-b-t and the other uses the path s-b-a-t, resulting in a total cost of $4x + y$. Therefore, $PoA_{sc}(\Delta) = \frac{4x+y}{4x}$, which can be arbitrarily high. □

Our goal is to characterize network topologies in which such a "bad" example cannot occur, i.e., topologies in which the PoA is always bounded, independently of the specific edge costs and capacities. The lower bound of n for a network with two parallel links motivates the following definition.

Definition 3 A symmetric network $G = (V, E)$ with source s and sink t is PoA-bounded for a family of symmetric CCS games $\mathcal{F}$ if for every symmetric CCS game $\Delta \in \mathcal{F}$ on the symmetric network G, it holds that $PoA(\Delta) \leq n$.

Our main result is a full characterization of PoA-bounded network topologies.

Theorem 4 *For symmetric CCS games, a symmetric network topology G is PoA-bounded with respect to sum-cost if and only if G is a series-parallel (SP) network.*

The proof of our characterization is composed of two parts. First, we show that for every symmetric CCS game that is played on a SP network $PoA_{sc} \leq n$. This is the content of Theorem 5. Second, we show that for every symmetric network topology G that is not a SP network, there exists a game that is played on G for which the PoA can be arbitrarily high. This part is the content of Theorem 7.

Theorem 5 *Let Δ be a feasible CCS game with an underlying graph G. If G is a SP graph then $PoA_{sc}(\Delta) \leq n$.*

Before presenting the proof, we establish the following lemma, which is used more than once throughout the paper.

Lemma 6 *Let Δ be a CCS game with an underlying SP graph G. Let S be a feasible profile of k agents and let S' be a feasible profile of r agents such that $r < k$. There exists a feasible s-t path in G that uses only edges that are used in the profile S that is feasible in S'.*

Proof The SP graph G is constructed by a sequence of series compositions and parallel compositions of SP graphs. This sequence can be viewed as a binary tree, in which every leaf is a single edge, and every inner node is either a series-composition operator or a parallel-composition operator. Given a SP graph G, this tree is termed the *construction tree* of G. The lemma is proved using an induction on the height of the construction tree.

The induction case is the case of a single edge, from which the assertion follows trivially. The induction hypothesis states the following: given a SP graph G whose construction tree is of height at most m, and strategy profiles S and S' as described above, given S', there exists a feasible s-t path in G that uses only edges that are used in S.

Assume the assertion holds for any tree of size m, and consider a SP graph with construction tree of height $m+1$, composing graphs G_1 and G_2. Let S_i and S_i' denote the induced profiles of S and S', respectively, played on graph G_i for $i \in 1, 2$.

We distinguish between two cases as follows:

Case (a): $G = G_1 || G_2$. Suppose that in S there are k' agents using G_1 (and $k - k'$ agents using G_2), and in S' there are r' agents using G_1 (and thus $r - r'$ agents using G_2). Since $r < k$, it holds that either $r' < k'$ or $r - r' < k - k'$; w.l.o.g. assume that $r' < k'$. By the induction hypothesis there exists a feasible s-t path in G_1 using only edges that are used in S_1. This path is also a feasible s-t path in G.

Case (b): $G = G_1 \to G_2$. Let s_i and t_i be the respective source and sink nodes of network G_i for $i \in 1, 2$. By the induction hypothesis, there exists a feasible s_i-t_i path using only edges that are used in S_i, for $i = 1, 2$. A feasible s-t path in G that uses only edges in S is obtained by concatenating these two paths. The assertion of the lemma follows. □

With this at hand, we are ready to prove Theorem 5.

Proof Given a game Δ, let S be a NE and S^* be an optimal profile with respect to sum-cost. We claim that the cost of every agent in S is at most the sum-cost of S^*. Assume by contradiction that there exists an agent i in the profile S whose cost is higher than $sc(S^*)$. By Lemma 6, given the profile S_{-i}, there exists a feasible s-t path that uses only edges that are used in S^*. The cost of this path is at most $sc(S^*)$; therefore, agent i benefits by deviating to this path, in contradiction to S being a

NE. It follows that the cost of every agent in S is at most $sc(S^*)$; therefore, the sum of the agents' costs in S cannot exceed $n \cdot sc(S^*)$. The assertion of the theorem follows. □

In order to complete the characterization it remains to show that for every non-SP network G, there exists a symmetric CCS game on G that has an unbounded price of anarchy.

Theorem 7 *Let G be a non-SP symmetric network. Then, there exists a symmetric CCS game on G for which the price of anarchy is arbitrarily high.*

In order to prove the last theorem, we use the following result, established by Milchtaich [25].

Lemma 8 [25] *A symmetric network G is a SP network if and only if the symmetric network in* Fig. 2 *is not embedded in G.*

The network topology in the last lemma is precisely the network topology with the unbounded PoA that motivated our study (see Proposition 2). The last lemma asserts that this graph topology is embedded in every non-SP network. Thus, in order to establish the assertion of Theorem 7, it remains to show that the unbounded PoA given in Proposition 2 can be *extended* to every network topology that embeds it. This is established in the following lemma.

Lemma 9 *Let G be a symmetric network that is not PoA-bounded with respect to sum-cost for a family of symmetric CCS games $\mathcal{F}$, and suppose G is embedded in a*

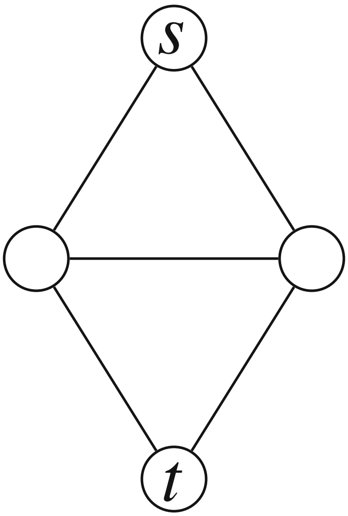

Fig. 2 A Braess Graph

symmetric network G'. Then, G' is not PoA-bounded with respect to sum-cost for the family $\mathcal{F}$ either.

Proof For every symmetric CCS game on the network G, there exists an equivalent symmetric CCS on the network G', which is obtained by setting edge costs and capacities as follows: for a subdivision operation of an edge e into two edges e_1 and e_2, assign edge costs $p_{e_1} = 0$ and $p_{e_2} = p_e$ and capacities $c_{e_1} = c_{e_2} = c_e$. For an addition operation of an edge e, set $c_e = 0$. For an extension operation of an edge e, set $p_e = 0$ and $c_e = n$. The original game (played on G) can be simulated in the network G' using the appropriate capacities and costs as specified above. It is easy to verify that if the embedded network G is not PoA-bounded with respect to sum-cost for a given CCS game, then this is also the case for any other network G', such that G is embedded in G'. □

This established the assertion of Theorem 7.

3.1.1 Parallel-Edge Networks

For the case of parallel-edge networks, we show that the PoA cannot exceed the maximum edge capacity in the network.

Theorem 10 *Let Δ be a feasible CCS game with an underlying graph G that consists of parallel edges. Let C_m denote the maximum capacity of any edge in G. It holds that $PoA_{sc}(\Delta) \leq C_m$.*

Proof Let S and S^* be a worst NE and an optimal solution in Δ, respectively. Denote by L the set of edges used only in S, by Q the set of edges used only in S^* and by M the set of edges used in both solutions, S and S^*. In addition, let x_e (respectively, x_e^*) denote the number of agents using edge e in profile S (resp., S^*). Also, for every set of edges $T \in \{L, Q, M\}$, let $X_T = \sum_{e \in T} x_e$ denote the total number of agents that use edges in T in the profile S, and let $X_T^* = \sum_{e \in T} x_e^*$ denote the total number of agents that use edges in T in the profile S^*. Similarly, let $P_T = \sum_{e \in T} p_e$ denote the total cost of edges in T.

We are now ready to state the proof of the theorem. Let $e' = argmin_{e \in Q} p_e$ be the cheapest edge in Q. By the definition of a NE, it holds that for every edge $e \in L$,

$$\frac{p_e}{x_e} \leq p_{e'}. \tag{2}$$

We distinguish between two cases, namely, the case in which $X_L \leq X_Q^*$ and the complementary case.

Case (a): $X_L \leq X_Q^*$. We get:

$$\sum_{e \in L} p_e = \sum_{e \in L} x_e \frac{p_e}{x_e} \leq p_{e'} \sum_{e \in L} x_e = p_{e'} X_L \leq p_{e'} X_Q^*,$$

where the first inequality follows from (2), and the last inequality follows by the assumption of case (a). On the other hand, by the definition of the edge $e' \in Q$, it follows that

$$\sum_{q \in Q} p_q \geq \lceil \frac{x^*_Q}{C_m} \rceil \cdot p_{e'}.$$

Therefore, $PoA(\Delta) \leq \frac{x^*_Q \cdot p_{e'}}{\lceil \frac{X^*_Q}{C_m} \rceil \cdot p_{e'}} \leq C_m$.

Case (b): $X_L > X^*_Q$. In this case we further distinguish between the case where $Q = \emptyset$ and the case where $Q \neq \emptyset$.

We begin with the case where $Q = \emptyset$. One can easily verify that in every NE, there is at most one edge that is neither empty nor full. To see this, note that if there are two partially full edges, simple arithmetics show that the fact that no agent wishes to deviate in one direction necessarily implies that there exists an agent that wishes to deviate in the opposite direction. Since $X_L > X^*_Q$, the partially-full edge must belong to M; denote it by $\hat{e}$. This further implies that $X_L \leq C_m - 1$.

By the definition of $\hat{e}$, for every $e \in L$, $p_e/c_e \leq p_{\hat{e}}/2$. Therefore, $P_L \leq \frac{1}{2} p_{\hat{e}} X_L \leq \frac{1}{2} p_{\hat{e}}(C_m - 1)$, where the last inequality follows from the observation above that $X_L \leq C_m - 1$.

On the other hand, it clearly holds that $P_M \geq p_{\hat{e}}$, as $\hat{e} \in M$.

We get that

$$PoA(\Delta) = \frac{P_L + P_M}{P_M} = 1 + \frac{P_L}{P_M} \leq 1 + \frac{\frac{1}{2} p_{\hat{e}}(C_m - 1)}{p_{\hat{e}}} = 1 + \frac{1}{2}(C_m - 1) \leq C_m,$$

where the last inequality holds for every $C_m \geq 1$.

We now prove the case in which $Q \neq \emptyset$. Recall that we are in case (b), where $X_L > X^*_Q$. Let $\lambda = X_L - X^*_Q$. Note that $1 \leq \lambda \leq C_m - 1$ due to the observation above that in every NE there could be at most one edge that is neither full nor empty. As before, let e' be the cheapest edge in Q.

To complete the proof, we use both (2) and the following observation. Let $\hat{e}$ be the partially full edge in M (as before). By the definition of NE, for every $e \in L$,

$$\frac{p_e}{x_e} \leq \frac{p_{\hat{e}}}{x_{\hat{e}} + 1}. \tag{3}$$

Recall that $X_L = X^*_Q + \lambda$. We can bound the individual cost of every agent in L in the profile S both by (2) and by (3). For X^*_Q of the agents, we will use (2), and for the remaining λ agents, we will use (3). We get $P_L \leq X^*_Q p_{e'} + \lambda \frac{p_{\hat{e}}}{x_{\hat{e}}+1}$.

Note that $P_Q \geq \frac{X^*_Q}{C_m} p_{e'}$; hence, $X^*_Q \cdot p_{e'} \leq C_m \cdot P_Q$. In addition, since $\lambda \leq C_m - 1$, and $x_{\hat{e}} \geq 1$, it follows that $\lambda \frac{p_{\hat{e}}}{x_{\hat{e}}+1} \leq (C_m - 1)\frac{p_{\hat{e}}}{2}$. Using the fact that $p_{\hat{e}} \leq P_M$, we get $P_L \leq C_m P_Q + (C_m - 1) P_M$, and therefore $P_L + P_M \leq C_m(P_Q + P_M)$.

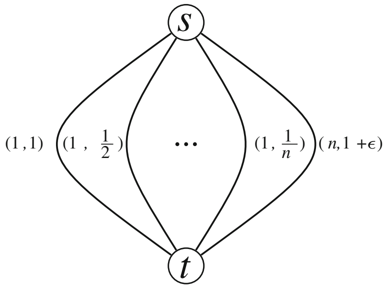

Fig. 3 The PoS with respect to sum-cost is $H(n)$

In conclusion,

$$PoA(\Delta) = \frac{P_L + P_M}{P_Q + P_M} \leq \frac{C_m(P_Q + P_M)}{P_Q + P_M} \leq C_m.$$

The assertion of the theorem follows. □

Theorems 5 and 10 assert that $PoA(\Delta) \leq \min\{n, C_m\}$. If the number of agents is big, i.e., n is fairly big compared to C_m, the effective PoA is bounded by C_m and therefore the result obtained in Theorem 10 is stronger.

3.2 Price of Stability (PoS)

In uncapacitated symmetric games, $PoS = 1$. In capacitated games, however, the PoS need not be optimal. Moreover, sub-optimality is obtained already in parallel-link networks.[2]

Theorem 11 *There exists a symmetric CCS game in which the PoS with respect to sum-cost is $H(n)$.*

Proof Consider a CCS game with n agents played on a graph that consists of $n+1$ parallel links, $e_1, \ldots, e_{n+1}$, such that for $i \in [n]$, $p_i = 1/i$ and $c_i = 1$; and $p_{n+1} = 1 + \epsilon$ and $c_{n+1} = n$ (see Fig. 3). It is easy to verify that the optimal solution is achieved when all the agents share edge e_{n+1}. However, this profile is not a NE since a single agent can benefit by deviating to edge e_n, thereby incurring a cost of $1/n$ instead of $(1+\epsilon)/n$. Following similar reasonings, agents will continue to deviate, one by one, until reaching the profile in which for every agent $i \in [n]$, agent i uses edge e_i. The cost of this profile is $H(n)$; the assertion follows. □

[2] A similar example is given in [32] for more general non-decreasing cost functions.

Using the potential-function method [4], it is established that the last bound is in fact tight. In particular, the proof uses the potential function $\phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} p_e/i$ and follows the same reasoning as in the uncapacitated case.

Theorem 12 *For every feasible symmetric CCS game, it holds that* $PoS_{sc} \leq H(n)$.

3.3 Homogeneous Capacities

It is interesting to note that in the special case in which all edges have the same capacity, it holds that $PoS_{sc} = 1$. To see this, observe that an optimal solution chooses $\lceil \frac{n}{C} \rceil$ disjoint s-t paths with minimum total cost. We claim that the profile in which the cheaper $\lfloor \frac{n}{C} \rfloor$ of these paths are saturated (and, possibly, the most expensive one is not full) is a NE. Indeed, a beneficial deviation will contradict the minimality of the total cost of the chosen paths. It follows that $PoS = 1$.

One might hope that homogeneity of capacities improves the PoA_{sc} as well. Unfortunately, the following example demonstrates that the PoA may still be unbounded. In particular, consider a game with two agents played on the graph depicted in Fig. 4. An optimal solution is one in which each agent uses a path that does not include the middle edge, for a total cost of 5. However, there exists an equilibrium in which one agent uses the path s-a-b-t (using the b-t edge of cost 1) and the other uses the path s-b-t (using the b-t edge of cost x). The total cost of this profile is $4 + x$, which can be arbitrarily high compared to 5.

4 The Max-Cost Objective Function

4.1 Price of Anarchy (PoA)

We first observe that the PoA can be arbitrarily high also with respect to the max-cost function.

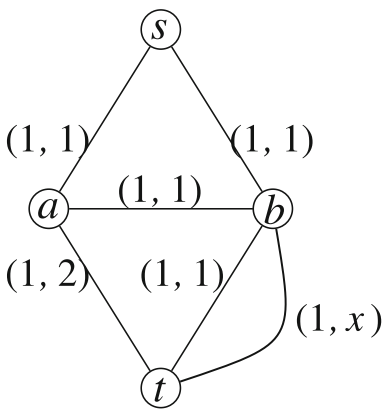

Fig. 4 The price of anarchy can be arbitrarily high, even with homogeneous capacities

Proposition 13 *The PoA with respect to max-cost in CCS games can be arbitrarily high.*

Proof As in the proof of Proposition 2, consider a game with two agents played on the graph depicted in Fig. 1. The optimal solution w.r.t. max-cost is the same as the one for sum-cost, and entails a cost of $2x$. However, in the worst NE, the maximal cost is $2x + y/2$. Therefore, $PoA_{mc} = \frac{2x+y/2}{2x}$, which can be arbitrarily high. □

As in the sum-cost case, we wish to characterize network topologies in which the PoA cannot exceed n. Interestingly, we obtain the exact same characterization as in the sum-cost case.

Theorem 14 *A symmetric network topology G is PoA-bounded w.r.t. max-cost if and only if G is a SP network.*

The necessity direction is established through similar arguments to the sum-cost function. In particular, by Proposition 13, the game given in Proposition 2 has an unbounded PoA with respect to the max-cost objective as well. This is exactly the topology given in 8, and it is easy to verify that Lemma 9 holds with respect to the max-cost objective function as well. Thus, the following theorem is established.

Theorem 15 *Let G be a non-SP symmetric network. Then, there exists a symmetric CCS game on G for which the price of anarchy is arbitrarily high.*

The sufficiency of SP networks is established in the following theorem.

Theorem 16 *Let Δ be a feasible CCS game with an underlying graph G. If G is a SP graph, then $PoA_{mc}(\Delta) \leq n$.*

Proof Let S be a NE and let S^* be an optimal solution with respect to max-cost. It follows by Lemma 6 that the cost of every agent in S is at most $sc(S^*)$. We get $mc(S) \leq sc(S^*) \leq n \cdot mc(S^*)$, where the second inequality follows since, by definition, for every profile T, $mc(T) \geq sc(T)/n$. The assertion of the theorem follows. □

4.1.1 Parallel-Edge Networks

For the case of parallel-edge networks, we show that the PoA cannot exceed the maximum edge capacity in the network.

Theorem 17 *Let Δ be a feasible CCS game with an underlying graph G that consists of a parallel edge. Let C_m denote the maximum cost of any edge in G. It holds that $PoA_{mc}(\Delta) \leq C_m$.*

Proof Let S and S^* be one of the worst NE and an optimal solution in Δ, respectively. Denote by L the set of edges used only in S, by Q the set of edges used only in S^* and by M the set of edges used in both solutions, S and S^*. In addition, let x_e (respectively, x_e^*) denote the number of agents using edge e in profile S (resp., S^*).

We distinguish between two cases, as follows:

Case (a): the maximum cost of any agent in the profile S^* is obtained on an edge $e \in Q$. Let e' be the cheapest edge in Q. Clearly, the max-cost in S^* is at least $p_{e'}/C_m$. On the other hand, by the definition of NE, no agent in S incurs a cost higher than $p_{e'}$ (otherwise, the agent that incurs a higher cost will deviate to e'). Therefore $PoA \leq C_m$.

Case (b): the maximum cost of any agent in the profile S^* is obtained on an edge $e \in M$. We further distinguish between two cases.

If the maximum cost of any agent in the profile S is also obtained on an edge in M, then let $e_1 \in M$ and $e_2 \in M$ denote the edges that realize the max cost in S and S^*, respectively. Since e_2 is more costly per agent in S^* than e_1, it holds that $\frac{p_{e_2}}{x^*_{e_2}} \geq \frac{p_{e_1}}{x^*_{e_1}}$. Since $x^*_{e_1} \leq C_m$, it follows that $p_{e_1} \leq C_m \frac{p_{e_2}}{x^*_{e_2}}$, which implies that $PoA \leq C_m$.

If the maximum cost of any agent in the profile S is obtained on an edge in L, then let $\ell \in L$ denote the edge that realizes the max cost in S.

If $Q = \emptyset$ and $S \neq S^*$, then there exists one edge in M that is neither full nor empty; denote this edge by $m \in M$. Since ℓ realizes the max cost in S, then $\frac{p_\ell}{x_\ell} \geq \frac{p_m}{x_m}$. On the other hand, by the definition of NE, it holds that $\frac{p_\ell}{x_\ell} \leq \frac{p_m}{x_m+1}$. A contradiction is reached.

Assume that $Q \neq \emptyset$, and let $m \in M$ denote the edge that realizes the max cost in S^*. Let e' denote the cheapest edge in Q. Since m is the edge that realizes the max cost in S^*, then $\frac{p_m}{x^*_m} \geq \frac{p_{e'}}{x^*_{e'}}$. Since $x^*_{e'} \leq C_m$, it follows that $p_{e'} \leq C_m \frac{p_m}{x^*_m}$. By the definition of NE, it holds that $\frac{p_\ell}{x_\ell} \leq p_{e'}$, and we get $\frac{p_\ell}{x_\ell} \leq C_m \frac{p_m}{x^*_m}$; i.e., $PoA \leq C_m$. □

4.2 Price of Stability (PoS)

For SP graphs, it follows directly from Theorem 16 that the PoS is bounded by n (since PoS is always bounded by PoA). This bound is tight, as follows from the example given in the proof of Theorem 11 (see Fig. 3). In this example, the unique NE is one in which every agent uses a distinct path, and the maximal cost incurred by any agent is 1, compared to $1/n$ in the optimal solution. For general networks, we establish the following bound.

Theorem 18 *For every CCS game Δ, it holds that $PoS_{mc}(\Delta)$ is bounded by $nH(n)$.*

Proof Consider the function $\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{x_e(S)}$. It is shown by [4] that this is an exact potential function for the game; i.e., it emulates the change in the cost of a deviating agent. It is easy to verify that for every profile T,

$$sc(T) \leq \Phi(T) \leq H(n) \cdot sc(T). \tag{4}$$

Let S^* be an optimal solution with respect to max-cost, and consider a NE S that is obtained by running best-response dynamics with an initial profile S^*. We get that $mc(S) \leq sc(S) \leq \Phi(S) \leq \Phi(S^*) \leq H(n)sc(S^*) \leq nH(n)mc(S^*)$, where the second and fourth inequalities follow from (4), the third inequality follows from the fact that Φ is a potential function and S is obtained from S^* through best-response steps, and the last inequality follows from the definition of max-cost. It follows that $mc(S)/mc(S^*) \leq H(n) \cdot n$, as promised. □

The bound presented above is not tight and the closing the gap between n and $n{\cdot}H(n)$ remains an open problem.

5 Convergence Rate of Best-Response Dynamics (BRD)

In this section we study the convergence rate of Best-Response Dynamics (BRD) to a NE. While BRD may in general be exponentially long [4], the following proposition establishes that in the case of a symmetric, undirected graph, BRD converges to a pure NE in at most n steps, and that this bound is tight. The intuition for this observation is that, in the uncapacitated version, after an agent deviates to some path P (as its best-response), the cost incurred by an agent using this path in the next iteration can only decrease; therefore, P remains a best-response move until all agents converge to the same path.

Observation 19 *For every uncapacitated cost-sharing game, every BRD converges to a NE in at most n steps, independently of the initial profile.*

In contrast, the following proposition shows that the convergence process of a capacitated game may be longer. In particular, we establish a lower bound of $\Omega(n^{3/2})$, even for parallel-link graphs.

Proposition 20 *There exists a symmetric CCS game and a best-response dynamics with a convergence time of* $\Omega(n^{3/2})$.

Before presenting the proof, we provide the intuition for the long convergence rate through an example. The $\Omega(n^{3/2})$ bound is obtained through BRD in which at every stage, the agent who incurs the lowest cost (among all agents who can deviate) is the one who deviates. The instance that realizes this bound is one in which there are $\sqrt{n}+2$ edges (assume for simplicity that n is a perfect square), where for $i = 1, \dots, \sqrt{n}+1$, edge e_i has capacity i, and the last edge has capacity n. Originally, all the agents reside on the last edge, which has a very high cost. We construct a cost vector of the remaining edges, such that whenever a new edge (among the first $\sqrt{n}+1$ edges) is "activated" (i.e., becomes non-empty), it will cause a chain of improvement steps that will end up in a profile where the edge preceding it is empty. The process that this edge undergoes when it is first activated repeats recursively. The cost functions that lead to such dynamics satisfy $p_i \leq p_{i+1}$ for every i, but the gap is sufficiently small so that agents prefer edges with a higher index if the cost is shared among more agents.

For example, consider the illustration given in Fig. 5. This figure demonstrates BRD that follows the process described above. In particular, we construct an instance with capacities as described above and cost functions (to be specified soon) that induce the following BRD process. Originally, all of the agents reside in the costly edge. Agent 1 migrates to edge 1, and agent 2 migrates to edge 2. After edge 2 has agent 2, agent 1 is better off joining agent 2 on edge 2. Edge 2 is now full (since it has a capacity of 2). Then, agent 3 migrates to edge 1, and agent 4 migrates to edge 3. Once edge 3 contains agent 4, agents 1 and 3 join agent 4 on edge 3. Agent 2 is now left alone on edge 2, so it migrates to edge 1, which is cheaper.

More generally, once a new edge i is activated, one agent from every edge $j < i$ deviates to the new edge i. As a result, the remaining $(i-2)$ agents on edge $i-1$ migrate to edges $j < i-1$, one to each edge. This process results in edge $i-1$ being empty. Thus, in the next iteration, edge $i-1$ is activated, which initiates the same process, recursively. The challenge is, therefore, to count the number of times a particular edge is activated, and the number of iterations each such activation entails.

We are now ready to state the proof of the proposition.

Proof Let G be a graph of $k+1$ parallel edges, denoted $e_1, \ldots, e_{k+1}$, and let $n = k(k+1)/2$. For every $i \in [k+1]$, let c_i and p_i denote the capacity and cost of edge e_i, respectively. Suppose that for every $i \in [k]$, $c_i = i$ and $p_i = 1 + (i-1)\epsilon$, and edge e_{k+1} is such that $c_{k+1} = n,\ p_{k+1} = n$ (i.e., edge $k+1$ has high capacity, but is very expensive relative to the other edges). We choose ϵ to be sufficiently small such that the following inequality is satisfied:

$$p_j/(i+1) < p_i/i \text{ for every } i < j. \tag{5}$$

Consider BRD in which in every step, the agent who incurs the minimum cost (among the agents that can benefit by a deviation) deviates. This dynamic has n phases, where the j'th phase starts at the first time the j'th agent deviates from edge e_{k+1} to one of the k edges. Consider phase $j+1$. Before the deviation of agent $j+1$, the other j agents are assigned to a subset of the k edges, such that the edges in use are all saturated (otherwise they could benefit from deviating and would incur a lower cost than agent $j+1$). Agent $j+1$ will "activate" a new edge e_i (for some i that shall be determined soon). Due to the edge cost structure (see Eq. 5), one agent from every edge e_1 to e_{i-1} (in that order) will deviate to the new edge e_i. Following the above migrations, due to the price structure (in particular $p_i < p_j$ for $i < j$), the agents from edge e_{i-1} will deviate to edges e_{i-2} to e_1 (one to each of these edges), and will fill them up.

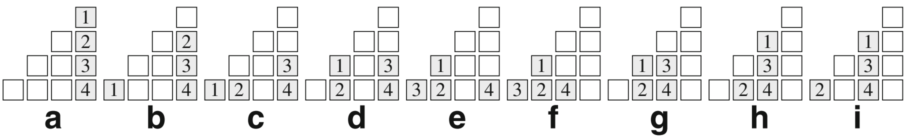

Fig. 5 Illustration of BRD with 4 agents

At this point, edge e_{i-1} is empty, and this will therefore be the edge to which the next agent from the expensive edge will migrate, causing the process to repeat itself.

In order to complete the analysis, we need to identify the edge to which agent $j + 1$ will deviate in its first deviation. Recall that this edge will be empty right before this deviation, and some subset of the other edges will be all saturated with the first j agents. Let m be the maximal number such that $j = m(m + 1)/2 + r$ for some $r \geq 0$. If $r = 0$, then it implies that edges e_1 to e_m are full (with a total capacity of $m(m + 1)/2$) and the $(j + 1)$'th agent will deviate to edge e_{m+1}. The process described above will follow, amounting to $2m$ steps during this phase. If $r > 0$, then among the first $m + 1$ edges, edge e_{m-r+1} is empty (the sum of the other m edges' capacities is exactly j). The deviation of agent $j + 1$ to e_{m-r+1} will be followed by $2m - 2r$ steps. At the end of this phase, edge e_{m-r+1} will be full and edge e_{m-r} will be empty and the same process will apply to edge e_{m-r} when agent $j + 2$ deviates.

It follows that for every $i \in [k]$, edge e_i is filled once during the first time it is activated, and once more for every edge e_ℓ such that $i < \ell \leq k$, i.e., $k - i + 1$ times. Each such occurrence takes $2i$ improvement steps. The total number of steps is, therefore, given by $\sum_{i=1}^{k}(k - i + 1)2i = \Omega(k^3) = \Omega(n^{3/2})$. The assertion of the theorem follows. □

6 Conclusions

In this work we introduce a model of capacitated network design games, and study the implications of edge capacities on the existence and quality of Nash equilibria with respect to different objective functions, as well as their implication on the convergence rate of best-response dynamics. We find that the consideration of edge capacities has a significant effect on all the above properties. Our main contribution is a full characterization of network topologies that have a bounded price of anarchy, independently of the edge capacities and costs.

Our results suggest many avenues for future research. A few obvious directions include studying asymmetric networks, better understanding the convergence rate of best-response dynamics, and closing the gap of the PoS with respect to the max-cost objective for general networks. In particular, in Section 4.2 we established an $n \cdot H(n)$ upper bound on the price of stability with respect to the max-cost objective for general graphs. We conjecture that the actual bound is n, and this remains an open question.

In addition, all the results presented here refer to single-commodity networks, where all the agents use the same sink and source. It will be interesting to study this problem in the more general case of multi-commodity networks (where each agent is associated with a different source and sink). Some of the results from the single-commodity case extend to the multi-commodity networks (such as the unbounded price of anarchy for non-series-parallel networks (under both the sum-cost and max-cost functions). Unfortunately, not all the results carry over.

In Section 5 we established a lower bound of $\Omega(n^{3/2})$ on the convergence rate of best-response dynamics in symmetric CCS games. Establishing an upper bound on the convergence rate of BRD in this case is an interesting open problem.

Finally, it will be interesting to consider additional objective functions.

Acknowledgments We thank Eli Ben-Sasson and Irit Dinur for helpful discussions.

References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. J. ACM **55**(6), 25:1–25:22 (2008)
2. Albers, S.: On the value of coordination in network design. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08, pp. 294–303. Society for Industrial and Applied Mathematics, Philadelphia (2008)
3. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: SODA'07 (2007)
4. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, pp. 295–304. IEEE Computer Society, Washington (2004)
5. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. In: STOC, pp. 511-520 (2003)
6. Bala, V., Goyal, S.: A noncooperative model of network formation. Econometrica **68**(5), 1181–1230 (2000)
7. Braess, D.: Uber ein paradoxon aus der verkehrsplanung. Unternehmensforschung **12**, 258–268 (1968)
8. Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC '05, pp. 99–107. ACM, New York (2005)
9. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Selfish routing in capacitated networks. Math. Oper. Res. **29**, 961–976 (2004)
10. Devanur, N.R., Mihail, M., Vazirani, V.V.: Strategyproof cost-sharing mechanisms for set cover and facility location games. In: Proceedings of ACM EC, pp. 108–114 (2003)
11. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. In: Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07, pp. 84–92. ACM, New York (2007)
12. Epstein, A., Feldman, M., Mansour, Y.: Efficient graph topologies in network routing games. Games Econ. Behav. **66**(1), 115–125 (2009)
13. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to nash equilibria. In: ICALP, pp. 502-513 (2003)
14. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, STOC '04, pp. 604–612. ACM, New York (2004)
15. Feldman, M., Tamir, T.: Convergence rate of best response dynamics in scheduling games with conflicting congestion effects. Working Paper (2011)
16. Fotakis, D.: Congestion games with linearly independent paths: convergence time and price of anarchy. Theory Comput. Syst. **47**(1), 113–136 (2010)
17. Hoefer, M., Skopalik, A.: Stability and convergence in selfish scheduling with altruistic agents. In: Proceedings of the 5th International Workshop on Internet and Network Economics, WINE '09, pp. 616–622. Springer-Verlag, Berlin, (2009)
18. Holzman, R., Nissan, L.-Y.: Strong equilibrium in congestion games. Games Econ. Behav. **21**(1-2), 85–101 (1997)
19. Holzman, R., Nissan, L.-y.: Network structure and strong equilibrium in route selection games. Math. Soc. Sci. **46**(2), 193–205 (2003)

20. Korilis, Y.A., Lazar, A.A., Orda, A.: Capacity allocation under noncooperative routing. IEEE Trans. Autom. Control **42**, 309–325 (1997)
21. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404–413 (1999)
22. Chen, H.l., Roughgarden, T.: Network design with weighted players. In: Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architextures (SPAA), pp. 29–38 (2006)
23. Milchtaich, I.: Topological conditions for uniqueness of equilibrium in networks. Math. Oper. Res. **30**, 225–244 (2005)
24. Milchtaich, I.: The equilibrium existence problem in finite network congestion games. In: Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE), pp. 87–98 (2006)
25. Milchtaich, I.: Network topology and the efficiency of equilibrium. Games Econ. Behav. **57**(2), 321–346 (2006)
26. Monderer, D.: Potential games. Games Econ. Behav. **14**(1), 124–143 (1996)
27. Papadimitriou, C.: Algorithms, games, and the internet. In: Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, STOC '01, pp. 749–753. ACM, New York (2001)
28. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. Int. J. Game Theory **2**(1), 65–67 (1973)
29. Roughgarden, T.: The price of anarchy is independent of the network topology. In: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing, STOC '02, pp. 428–437. ACM, New York (2002)
30. Roughgarden, T., Tardos, É.: How bad is selfish routing J. ACM **49**(2), 236–259 (2002)
31. Syrgkanis, V.: The complexity of equilibria in cost sharing games. In: Proceedings of the 6th international conference on Internet and network economics, WINE'10, pp. 366–377. Springer-Verlag, Berlin (2010)
32. von Falkenhausen, P., Harks, T.: Optimal cost sharing protocols for scheduling games. In: Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11, pp. 285–294. ACM, New York (2011)
33. Korilis, A.O.Y.A., Lazar, A.A.: Avoiding the braess paradox in noncooperative networks. J. Appl. Probab. **36**(1), 211–222 (1999)