

Congestion Games with Linearly Independent Paths: Convergence Time and Price of Anarchy

Dimitris Fotakis

Published online: 14 March 2009
© Springer Science+Business Media, LLC 2009

Abstract We investigate the effect of linear independence in the strategies of congestion games on the convergence time of best improvement sequences and on the pure Price of Anarchy. In particular, we consider symmetric congestion games on extension-parallel networks, an interesting class of networks with linearly independent paths, and establish two remarkable properties previously known only for parallel-link games. We show that for arbitrary (non-negative and non-decreasing) latency functions, any best improvement sequence reaches a pure Nash equilibrium in at most as many steps as the number of players, and that for latency functions in class \mathcal{D} , the pure Price of Anarchy is at most $\rho(\mathcal{D})$, i.e. it is bounded by the Price of Anarchy for non-atomic congestion games. As a by-product of our analysis, we obtain that for symmetric network congestion games with latency functions in class \mathcal{D} , the Price of Stability is at most $\rho(\mathcal{D})$.

Keywords Network congestion games · Best response dynamics · Price of anarchy · Price of stability

1 Introduction

Congestion games provide a natural model for non-cooperative resource allocation in large-scale communication networks and have been the subject of intensive research

A preliminary version of this work appeared in the Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT 2008), B. Monien and U.-P. Schroeder (Eds.), Lecture Notes in Computer Science 4997, pp. 33–45, Springer, 2008. This work was done while the author was with the Department of Information and Communication Systems Engineering, University of the Aegean, Greece.

D. Fotakis (✉)
School of Electrical and Computer Engineering, National Technical University of Athens,
15780 Athens, Greece
e-mail: fotakis@cs.ntua.gr

in algorithmic game theory. In a *congestion game*, a finite set of non-cooperative players, each controlling an unsplittable unit of load, compete over a finite set of resources. All players using a resource experience a latency (or cost) given by a non-negative and non-decreasing function of the resource's load (or congestion). Among a given set of resource subsets (or strategies), each player selects one selfishly trying to minimize her *individual cost*, that is the sum of the latencies on the resources in the chosen strategy. A natural solution concept is that of a *pure Nash equilibrium*, a configuration where no player can decrease her individual cost by unilaterally changing her strategy.

The prevailing questions in recent work on congestion games have to do with quantifying the inefficiency due to the players' selfish behaviour (see e.g. [4, 6, 7, 9, 10, 17, 22, 23]), and bounding the convergence time to pure Nash equilibria if the players select their strategies in a selfish and decentralized fashion (see e.g. [1, 8, 14, 21]). In this work, we investigate the effect of linear independence in the strategies of congestion games on the convergence time of best improvement sequences and on the inefficiency of pure Nash equilibria. In particular, we consider symmetric congestion games on extension-parallel networks, an interesting class of networks whose paths are linearly independent in the sense that every path contains an edge not included in any other path. For this class of congestion games, which comprises a natural and non-trivial generalization of the extensively studied class of parallel-link games (see e.g. [7, 14, 17, 21–23]), we provide best possible answers to both research questions above.

1.1 Convergence Time to Pure Nash Equilibria

In a seminal paper, Rosenthal [26] proved that the pure Nash equilibria of congestion games correspond to the local optima of a natural potential function. Hence Rosenthal established that every congestion game admits at least one pure Nash equilibrium reached in a natural way when players iteratively make best improvement moves, i.e. adopt strategies that minimize their individual cost given the strategies of other players. Nevertheless, this may take an exponential number of steps, since computing a pure Nash equilibrium is PLS-complete for non-symmetric network congestion games as shown by Fabrikant *et al.* [15]. In fact, the proof of Fabrikant *et al.* establishes the existence of instances where any sequence of players' best improvement moves is exponentially long. Even for symmetric network congestion games, where a pure Nash equilibrium can be found efficiently by a min-cost flow computation [15], Ackermann *et al.* [1] presented a class of instances where any best improvement sequence is exponentially long.

A natural approach to circumvent the strong negative results of [1, 15] is to resort to *approximate* pure Nash equilibria, where no player can *significantly* improve her individual cost by unilaterally changing her strategy. Chien and Sinclair [8] considered symmetric congestion games with a weak restriction on latency functions and proved that several natural families of best improvement sequences converge to an approximate equilibrium in polynomial time.

An orthogonal approach is to identify interesting classes of congestion games for which best improvement sequences reach a pure Nash equilibrium in a polynomial

number of steps. For instance, it is well known that for symmetric singleton congestion games (aka parallel-link games) with non-negative and non-decreasing latency functions, any best improvement sequence reaches a pure Nash equilibrium in at most n steps, where n denotes the number of players. Ieong *et al.* [21] proved that even for non-symmetric singleton games with non-monotonic latencies, best improvement sequences reach a pure Nash equilibrium in a polynomial number of steps. Subsequently, Ackermann *et al.* [1] generalized this result to *matroid* congestion games, where the strategy space of each player consists of the bases of a matroid over the set of resources. Furthermore, Ackermann *et al.* proved that the matroid property on the players' strategy spaces is necessary for guaranteeing polynomial convergence of best improvement sequences if one does not take into account the global structure of the game.

Contribution The negative results of [1, 15] leave open the possibility that some particular classes of symmetric network congestion games can guarantee fast convergence of best improvement sequences to pure Nash equilibria. In Sect. 3, we prove that for symmetric congestion games on extension-parallel networks with arbitrary (non-negative and non-decreasing) latency functions, any best improvement sequence reaches a pure Nash equilibrium in at most n steps¹ (Theorem 1). In particular, we show that in any best improvement sequence, every player moves at most once (Lemma 1). This result is best possible, since there are instances where reaching a pure Nash equilibrium requires that every player moves at least once. To the best of our knowledge, symmetric congestion games on extension-parallel networks is the only non-trivial generalization of parallel-link games for which such a strong bound on the convergence time of best improvement sequences is known.

1.2 Inefficiency of Pure Nash Equilibria

Having reached a pure Nash equilibrium, selfish players enjoy a minimum individual cost given the strategies of other players. However, the public benefit is usually measured by the *total cost* incurred by all players. Since a pure Nash equilibrium does not need to minimize the total cost, one seeks to quantify the inefficiency due to the players' non-cooperative and selfish behaviour. The *Price of Anarchy*, introduced by Koutsoupias and Papadimitriou [22], and the *Price of Stability*, introduced by Anshelevich *et al.* [5], are the prevailing measures of the performance degradation due to the players' non-cooperative and selfish behaviour. The (pure) Price of Anarchy (PoA) (resp. Price of Stability, PoS) is the *worst-case* (resp. *best-case*) ratio of the total cost of a (pure) Nash equilibrium to the optimal total cost.

¹We highlight that matroid games and games on extension-parallel networks have a different combinatorial structure and may have quite different properties. For example, a network consisting of two parallel-link networks connected in series is not extension-parallel, while the corresponding network congestion game is a symmetric matroid game. For another example, Milchtaich [25, Example 4] proved that weighted congestion games on extension-parallel networks may not admit a pure Nash equilibrium. On the other hand, Ackermann *et al.* [2, Theorem 2] proved that every weighted matroid congestion game admits at least one pure Nash equilibrium.

Many recent contributions have provided strong upper and lower bounds on the pure Price of Anarchy for several classes of congestion games, mostly congestion games with affine and polynomial latency functions and congestion games on parallel links.² Lücking *et al.* [23] and Suri *et al.* [29] were the first to consider the PoA of atomic congestion games for the objective of total cost. For (symmetric) parallel-link games with linear latency functions, Lücking *et al.* proved that the PoA is $4/3$. For non-symmetric singleton games with linear latency functions, Suri *et al.* proved that the PoA is at most $5/2$, which was later shown to be tight [7]. For parallel-link games with polynomial latency functions of degree d , Gairing *et al.* [17] proved the PoA is at most $d + 1$. Awerbuch *et al.* [6] and Christodoulou and Koutsoupias [9] proved independently that the PoA of congestion games is $5/2$ for affine latency functions and $d^{\Theta(d)}$ for polynomial latency functions of degree d . Subsequently, Aland *et al.* [4] obtained exact bounds on the PoA of congestion games with polynomial latency functions. In the non-atomic setting, where the number of players is infinite and each player controls an infinitesimal amount of load, Roughgarden [27] proved that the PoA is independent of the strategy space and equal to $\rho(\mathcal{D})$, where ρ depends on the class of latency functions \mathcal{D} only (e.g. ρ is equal to $4/3$ for affine and $\frac{27+6\sqrt{3}}{23}$ for quadratic functions). Subsequently, Correa *et al.* [11] introduced $\beta(\mathcal{D}) = 1 - \frac{1}{\rho(\mathcal{D})}$ and gave a simple proof of the same bound. Recently Fotakis [16] and independently Caragiannis *et al.* [7, Theorem 23] proved that the PoA of (atomic) parallel-link games with latency functions in class \mathcal{D} is at most $\rho(\mathcal{D})$, i.e. it is bounded by the PoA of non-atomic congestion games with arbitrary strategies.

There is no difference between the Price of Anarchy and the Price of Stability for non-atomic congestion games, since the Nash equilibrium in the non-atomic setting is essentially unique (under mild assumptions on the latency functions, see e.g. [28]). In the atomic setting, Christodoulou and Koutsoupias [10] proved that the PoS of congestion games with affine latencies lies between $1 + \sqrt{3}/3$ and 1.6 . Subsequently, Caragiannis *et al.* [7, Theorem 6] proved that the PoS of affine congestion games is $1 + \sqrt{3}/3$, and that for non-symmetric singleton games with latency functions in class \mathcal{D} , the PoS is at most $\rho(\mathcal{D})$.

Contribution Despite the considerable interest in the PoA and the PoS of congestion games, it remains open whether some better upper bounds close to ρ are possible for symmetric games on simple networks other than parallel links (e.g. extension-parallel networks, series-parallel networks), or strong lower bounds similar to the lower bounds of [4, 6, 9, 10] also apply to them. As a first step in this direction, we prove in Sect. 4 that the PoA of symmetric congestion games on extension-parallel networks with latency functions in class \mathcal{D} is at most $\rho(\mathcal{D})$ (Theorem 3). On the negative side, we show that this result cannot be further generalized to series-parallel networks. In particular, we present an example of a 3-player symmetric game on a simple series-parallel network with linear latencies for which the PoA is $15/11$,

²Here we cite only the most relevant results on the pure PoA for the objective of total cost. For a survey on the PoA of congestion games for the objective of total cost and the objective of maximum cost, see e.g. [18].

greater than the value of ρ for linear latencies. To the best of our knowledge, symmetric games on extension-parallel networks is the only non-trivial generalization of parallel-link games for which such a strong bound on the pure Price of Anarchy is known.

To establish the upper bound of $\rho(\mathcal{D})$ on the PoA of symmetric games on extension-parallel networks, we show two properties of symmetric games that may be of independent interest. In particular, we prove that (i) for symmetric games on general networks with latencies in class \mathcal{D} , the total cost of any configuration minimizing Rosenthal's potential function is at most $\rho(\mathcal{D})$ times the optimal total cost (Lemmas 2 and 3, see also [5, Theorem 3.4]), and that (ii) for symmetric games on extension-parallel networks, every pure Nash equilibrium is a global minimum of Rosenthal's potential function (Lemma 4, also implicit in the work of Holzman and Law-Yone [19, Theorem 6.1]).

The first property implies an upper bound of $\rho(\mathcal{D})$ on the PoS of symmetric network congestion games with latency functions in class \mathcal{D} (Theorem 2), which was also obtained by Anshelevich *et al.* [5, Theorem 3.4].³ For instance, this implies that the PoS of symmetric network games with affine latencies is at most $4/3$, i.e. better than the bound of $1 + \sqrt{3}/3$ known for general congestion games.

1.3 Related Work on Congestion Games with Linearly Independent Strategies

There has been a significant volume of previous work investigating the impact of linearly independent strategies on properties of congestion games. In [19], Holzman and Law-Yone proved that a strategy space admits a strong equilibrium⁴ for any selection of (non-negative and non-decreasing) latency functions if and only if it does not contain a so-called *bad configuration*, which is defined as a triplet of strategies s_1, s_2, s_3 such that s_1 includes some resource e_1 not belonging to s_2 , s_2 includes some resource e_2 not belonging to s_1 , and s_3 includes both resources e_1 and e_2 . Furthermore, they proved that in a strategy space without a bad configuration, every strategy s contains a resource not included in any other strategy that is not a superset of s . Holzman and Law-Yone showed that for symmetric congestion games without a bad configuration, every pure Nash equilibrium is a strong equilibrium. Subsequently, Holzman and Law-Yone [20] proved that the class of symmetric congestion games on extension-parallel networks is the network equivalent of symmetric congestion games without a bad configuration.

Milchtaich [24] was the first to consider networks with linearly independent paths (under this name). Milchtaich proved that an undirected network has linearly independent paths if and only if it is extension-parallel, and if and only if it does not contain a triplet of paths comprising a bad configuration. Moreover, Milchtaich showed that networks with linearly independent paths is the only class of networks where for any selection of non-negative and increasing (resp. non-decreasing) latency functions, all equilibria in the non-atomic setting are Pareto (resp. weakly Pareto) efficient.

³We note that our proof was obtained independently and is different from that of [5, Theorem 3.4].

⁴The notion of *strong equilibrium* was introduced by Aumann in late 50's. A configuration is a strong equilibrium if no coalition of players can deviate in a way profitable for all its members.

Recently Epstein *et al.* [12, 13] considered fair connection games⁵ and congestion games on extension-parallel networks. In [13], Epstein *et al.* proved that fair connection games on extension-parallel networks admit a strong equilibrium. In [12], Epstein *et al.* showed that extension-parallel networks is the only class of networks where for any selection of (non-negative and non-decreasing) latency functions, every pure Nash equilibrium minimizes the maximum players’ cost.

2 Model and Preliminaries

For any integer $k \geq 1$, we denote $[k] \equiv \{1, \dots, k\}$. For a vector $x = (x_1, \dots, x_n)$, we denote $x_{-i} \equiv (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i}, x'_i) \equiv (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

2.1 Congestion Games

A *congestion game* is a tuple $\Gamma(N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E})$, where N denotes the set of players, E denotes the set of resources, $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$ denotes the strategy space of each player i , and $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ is a non-negative and non-decreasing latency function associated with each resource e . A congestion game is *symmetric* if all players have a common strategy space. In the following, we let n denote the number of players.

A *configuration* is a vector $\sigma = (\sigma_1, \dots, \sigma_n)$ consisting of a strategy $\sigma_i \in \Sigma_i$ for each player i . For every resource e , we let $\sigma_e = |\{i \in N : e \in \sigma_i\}|$ denote the congestion induced on e by σ . The *individual cost* of player i in the configuration σ is $c_i(\sigma) = \sum_{e \in \sigma_i} d_e(\sigma_e)$. A configuration σ is a *pure Nash equilibrium* if no player can improve her individual cost by unilaterally changing her strategy. Formally, σ is a Nash equilibrium if for every player i and every strategy $s_i \in \Sigma_i$, $c_i(\sigma) \leq c_i(\sigma_{-i}, s_i)$.

Rosenthal [26] introduced the potential function

$$\Phi(\sigma) = \sum_{e \in E} \sum_{i=1}^{\sigma_e} d_e(i),$$

and proved that when a player i switches from her current strategy σ_i to a new strategy $s_i \in \Sigma_i$, the difference in the potential value equals the difference in the individual cost of player i , i.e.

$$\Phi(\sigma_{-i}, s_i) - \Phi(\sigma) = c_i(\sigma_{-i}, s_i) - c_i(\sigma)$$

Thus Rosenthal established that congestion games (even with non-monotonic latency functions) admit a pure Nash equilibrium, and that the pure Nash equilibria of a congestion game correspond to the local minima of the potential function Φ . For symmetric network congestion games with arbitrary (non-negative and non-decreasing)

⁵In fair connection games, each player has a source and a sink, and connects them by selecting a path in the underlying network. The edge costs are shared equally among the players using them. Anshelevich *et al.* [5, Theorem 2.1] observed that fair connection games fall into the class of congestion games as defined by Rosenthal [26]. Hence fair connection games admit a potential function and have at least one pure Nash equilibrium.

latency functions, Fabrikant *et al.* [15, Theorem 2] proved that the global minimum of the potential function Φ , and thus a pure Nash equilibrium, can be computed in polynomial time by a min-cost flow computation.

In this work, we focus on *symmetric network congestion games*, where the players' strategies are determined by a directed network $G(V, E)$ with a distinguished source s and sink t (aka a $s - t$ network). The network edges play the role of resources and the common strategy space of the players is the set of (simple) $s - t$ paths in G , denoted \mathcal{P} . For any $s - t$ path p and any pair of vertices v_1, v_2 on p , we let $p[v_1, v_2]$ denote the segment of p between vertices v_1 and v_2 ($p[v_1, v_2]$ is empty if v_1 appears after v_2 on p). For consistency with the definition of strategies as resource subsets, we usually regard paths as sets of edges.

2.2 Flows and Configurations

Let $G(V, E)$ be a directed network with source s and sink t . A $s - t$ flow f is a vector $(f_e)_{e \in E} \in \mathbb{R}_{\geq 0}^m$ that satisfies the flow conservation at all vertices other than s and t . The volume of f is the total flow leaving s . A flow f is *acyclic* if there is no directed cycle in G with positive flow on all its edges. For a flow f and a path $p \in \mathcal{P}$, we let $f_p^{\min} \equiv \min_{e \in p} \{f_e\}$.

Given a configuration σ of a symmetric network congestion game Γ , we refer to the congestion vector $(\sigma_e)_{e \in E}$ as the $s - t$ flow induced by σ . We say that a flow f is *feasible* if there is a configuration inducing congestion f_e on every edge e . Hence any configuration of Γ corresponds to a feasible flow, while a feasible flow may be induced by many different configurations. Moreover, every integer acyclic $s - t$ flow of volume n corresponds to (possibly many) configurations of Γ , and thus is feasible. We say that a configuration is *acyclic* if the corresponding feasible flow is acyclic. In the following, we slightly abuse the notation by letting the same symbol denote both a configuration and the feasible flow induced by it.

2.3 Best Improvement Sequences

A strategy $s_i \in \Sigma_i$ is a *best response* of player i to a configuration σ (or equivalently to σ_{-i}) if for every strategy $s'_i \in \Sigma_i$, $c_i(\sigma_{-i}, s_i) \leq c_i(\sigma_{-i}, s'_i)$. If the current strategy σ_i of player i is not a best response to the current configuration σ , a best response of i to σ is called a *best improvement* of player i . We consider *best improvement sequences*, where in each step, a player i whose strategy σ_i is not a best response to the current configuration σ switches to her best improvement. The existence of the potential function Φ implies that any best improvement sequence converges to a pure Nash equilibrium in a finite number of steps. In the first part of this work, we are interested in bounding the maximum length of best improvement sequences for symmetric congestion games on extension-parallel networks.

2.4 Social Cost, Price of Anarchy, and Price of Stability

In the second part of this work, we are interested in quantifying the inefficiency of pure Nash equilibria for symmetric congestion games on extension-parallel networks

and on general networks. We evaluate configurations using the objective of *total cost*. The total cost $C(\sigma)$ of a configuration σ is the sum of players' individual costs in σ , i.e.

$$C(\sigma) = \sum_{i=1}^n c_i(\sigma) = \sum_{e \in E} \sigma_e d_e(\sigma_e)$$

The *optimal configuration*, denoted o , minimizes the total cost among all configurations in \mathcal{P}^n .

The pure *Price of Anarchy* (PoA) of a congestion game Γ is the maximum ratio $C(\sigma)/C(o)$ over all pure Nash equilibria σ of Γ . The pure *Price of Stability* (PoS) of Γ is the minimum ratio $C(\sigma)/C(o)$ over all pure Nash equilibria σ of Γ . In other words, the pure PoA (resp. PoS) is equal to $C(\sigma)/C(o)$, where σ is Γ 's pure Nash equilibrium of maximum (resp. minimum) total cost. Hence the pure PoA (resp. PoS) of Γ is bounded from above by a number α if *all* pure Nash equilibria (resp. *some* pure Nash equilibrium) of Γ have (resp. has) total cost at most α times the optimal total cost.

We use the quantities $\rho(\mathcal{D})$ and $\beta(\mathcal{D})$ introduced in [11, 27] respectively, to bound the PoS of symmetric network congestion games and the PoA of symmetric games on extension-parallel networks. For a non-negative and non-decreasing function $d(x)$, let $\rho(d) \equiv \sup_{x \geq y \geq 0} \frac{xd(x)}{yd(y)+(x-y)d(x)}$, and let $\beta(d) \equiv \sup_{x \geq y \geq 0} \frac{y(d(x)-d(y))}{xd(x)}$. For a non-empty class of non-negative and non-decreasing functions \mathcal{D} , let $\rho(\mathcal{D}) \equiv \sup_{d \in \mathcal{D}} \rho(d)$, and let $\beta(\mathcal{D}) \equiv \sup_{d \in \mathcal{D}} \beta(d)$. We note that $(1 - \beta(\mathcal{D}))^{-1} = \rho(\mathcal{D})$.

2.5 Extension-Parallel Networks

Let $G_1(V_1, E_1)$ and $G_2(V_2, E_2)$ be two networks with source $s_1 \in V_1$ and $s_2 \in V_2$ and sink $t_1 \in V_1$ and $t_2 \in V_2$ respectively, and let $G'(V_1 \cup V_2, E_1 \cup E_2)$ be the union network of G_1 and G_2 . The *parallel composition* of G_1 and G_2 results in a $s - t$ network obtained from G' by identifying s_1 and s_2 to the source s and identifying t_1 and t_2 to the sink t . The *series composition* of G_1 and G_2 results in a $s - t$ network obtained from G' by letting s_1 be the source s , letting t_2 be the sink t , and identifying t_1 with s_2 .

A directed $s - t$ network is *series-parallel* if it consists of either a single edge (s, t) or two series-parallel networks composed either in series or in parallel. A directed $s - t$ network is *extension-parallel* if it consists of either: (i) a single edge (s, t) , (ii) a single edge and an extension-parallel network composed in series, or (iii) two extension-parallel networks composed in parallel. Every extension-parallel network is series-parallel. The converse is true only if in every series composition, at least one component is a single edge.

The following proposition gives a useful property of series-parallel (and thus of extension-parallel) networks. For completeness, we prove the proposition in the [Appendix](#), Sect. A.1.

Proposition 1 *Let $G(V, E)$ be a series-parallel $s - t$ network, and let u, w be two vertices connected by two internally disjoint paths p and p' . Then every $s - t$ path having an edge in common with p contains both u and w and does not intersect p' at any vertex other than u and w .*

A $s - t$ network has *linearly independent paths* if every $s - t$ path contains at least one edge not belonging to any other $s - t$ path.⁶ Milchtaich [24, Proposition 5] proved that an undirected $s - t$ network has linearly independent paths if and only if it is extension-parallel. Therefore, every (directed) extension-parallel network has linearly independent paths (see also [20, Theorem 1]). Furthermore, [24, Propositions 3 and 5] imply that a (directed) series-parallel network has linearly independent paths if and only if it is extension-parallel.

The following proposition gives a useful property of extension-parallel networks (for a proof, see [24, Proposition 4]).

Proposition 2 (Milchtaich) *Let G be an extension-parallel $s - t$ network. Then for every pair of $s - t$ paths p and p' in G , and every vertex v common to both paths, p and p' share either the segment between s and v or the segment between v and t .*

Proposition 2 implies that for any two paths p and p' of an extension-parallel network, the segments $p \setminus p'$ and $p' \setminus p$ where p and p' deviate from each other are two internally disjoint paths with common endpoints.

The following proposition gives an interesting property of networks with linearly independent paths (and thus of extension-parallel networks).

Proposition 3 *Let Γ be a symmetric congestion game on a $s - t$ network G with linearly independent paths, let f be any configuration of Γ , and let π be any (simple) path of G with $f_\pi^{\min} > 0$. Then there exists a player i whose strategy in f includes π .*

Proof We first consider the case where π is a $s - t$ path. Since G is a $s - t$ network with linearly independent paths, there exists an edge $e \in \pi$ not belonging to any other $s - t$ path. Let i be any player whose strategy in f includes e (such a player exists because $f_e > 0$). Then the strategy of player i in f is π (since e would not be included in f_i otherwise).

If π is not a $s - t$ path, let u and w be the endpoints of π . Since f corresponds to a $s - t$ flow with $f_\pi^{\min} > 0$, there is a $s - u$ path π_1 with $f_{\pi_1}^{\min} > 0$ (π_1 is empty if u is s) and a $w - t$ path π_2 with $f_{\pi_2}^{\min} > 0$ (π_2 is empty if w is t). Then $p = (\pi_1, \pi, \pi_2)$ is a $s - t$ path with $f_p^{\min} > 0$. As shown before, there is a player i whose strategy in f is p , which includes π . \square

Every configuration of a symmetric congestion game on a series-parallel (and thus on an extension-parallel) network corresponds to a (feasible) acyclic flow of volume n . Proposition 3 implies that for every symmetric congestion game Γ on an extension-parallel $s - t$ network, any acyclic $s - t$ flow of volume n corresponds to a unique configuration of Γ (uniqueness is up to players' permutation). Therefore, for symmetric congestion games on extension-parallel networks, there is a correspondence between configurations and feasible acyclic flows.

⁶The name is motivated by the fact that in such a network, it is not possible to express any path as the symmetric difference of some other paths [24, Proposition 6].

3 Convergence Time to Pure Nash Equilibria

In this section, we prove that for symmetric congestion games on extension-parallel networks, any best improvement sequence reaches a pure Nash equilibrium in at most n steps. The proof is based on the following lemma establishing that in any best improvement sequence, every player moves at most once.

Lemma 1 *Let Γ be a symmetric congestion game on an extension-parallel network, let σ be the current configuration, and let i be a player switching from her current strategy σ_i to her best improvement σ'_i . Then for every player j whose current strategy σ_j is a best response to σ , σ_j remains a best response of j to the new configuration $\sigma' = (\sigma_{-i}, \sigma'_i)$.*

Proof For sake of contradiction, we assume that there is a player j whose current strategy σ_j is a best response to σ but not to σ' . Let σ'_j be the best response of j to σ' , and let $p = \sigma_j \setminus \sigma'_j$ and $p' = \sigma'_j \setminus \sigma_j$ be the segments where σ_j and σ'_j deviate from each other. Due to the extension-parallel structure of the network (see also Proposition 2), p and p' are internally disjoint paths with common endpoints, denoted u and w (u may be s and w may be t).

Since p and p' are edge-disjoint paths and player j improves her individual cost in σ' by switching from p to p' ,

$$\sum_{e \in p} d_e(\sigma'_e) > \sum_{e \in p'} d_e(\sigma'_e + 1) \tag{1}$$

Using the inequality above and the hypothesis that σ'_i is a best improvement of player i to σ , and exploiting the extension-parallel structure of the network, we establish that if player j prefers σ'_j to σ_j in the new configuration σ' , then σ_j is not a best response of j to σ . In particular, we show that player j can also improve her individual cost in σ by switching from an appropriate segment of σ_j to the corresponding segment of σ'_i . The precise structure of these segments depends on the subgraph defined by σ_j , σ'_j , σ_i , and σ'_i , and dictates the different cases considered in the proof (see also Fig. 1). In all cases, we obtain a contradiction to the hypothesis that σ_j is a best response of j to σ .

The technical part of the proof proceeds by case analysis. At a high level, we distinguish between two cases depending on whether σ'_i contains both u and w , in which case $\sigma'_i[u, w]$ can serve as an alternative to p , or not.

Case I: $u, w \in \sigma'_i$ We first consider the case where σ'_i contains u and w . We further distinguish between two subcases depending on whether p and σ'_i are edge-disjoint or not.

Case Ia: $p \cap \sigma'_i = \emptyset$ We start with the case where p and σ'_i are edge-disjoint. We first consider the case where $\sigma'_i[u, w] \setminus p'$ does not contain any edges of σ_i (see e.g.

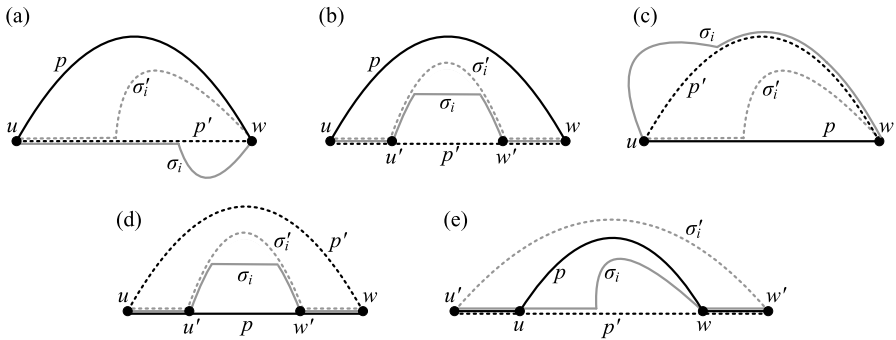


Fig. 1 The different cases considered in the proof of Lemma 1. In each case, the *solid black path* labeled p represents the best response of player j in σ between vertices u and w , the *solid grey path* labeled σ_i represents the strategy of player i in σ , and the *dotted grey path* labeled σ'_i represents the best response of player i to σ . We assume that the best response of player j changes from p to the *dotted black path* labeled p' when player i switches from σ_i to σ'_i , and establish a contradiction in all cases

Fig. 1a). Then,

$$\begin{aligned}
 \sum_{e \in p'} d_e(\sigma'_e + 1) &\geq \sum_{e \in p' \cap \sigma'_i} d_e(\sigma_e + 1) + \overbrace{\sum_{e \in (p' \cap \sigma_i) \setminus \sigma'_i} d_e(\sigma_e) + \sum_{e \in (p' \setminus \sigma_i) \setminus \sigma'_i} d_e(\sigma_e + 1)}^{\text{cost of player } i \text{ on } p' \setminus \sigma'_i \text{ in } \sigma} \\
 &\geq \sum_{e \in p' \cap \sigma'_i} d_e(\sigma_e + 1) + \underbrace{\sum_{e \in \sigma'_i[u, w] \setminus p'} d_e(\sigma_e + 1)}_{\text{cost of player } i \text{ on } \sigma'_i[u, w] \setminus p' \text{ in } \sigma} \\
 &= \sum_{e \in \sigma'_i[u, w]} d_e(\sigma_e + 1) \tag{2}
 \end{aligned}$$

For the first inequality, we use that when player i switches from σ_i to σ'_i : (i) the congestion of any edge e in σ'_i does not decrease (i.e. $\forall e \in \sigma'_i, \sigma'_e \geq \sigma_e$), (ii) the congestion of any edge e decreases by at most 1 (i.e. $\forall e, \sigma'_e \geq \sigma_e - 1$), and (iii) the congestion of any edge e not in σ_i or σ'_i does not change (i.e. $\forall e \notin \sigma_i \cup \sigma'_i, \sigma'_e = \sigma_e$). For the second inequality, we observe that the marked terms on the left-hand side is the individual cost of player i on $p' \setminus \sigma'_i$ in σ (i.e. when the configuration of the remaining players is σ_{-i}) and that the marked term on the right-hand side is the individual cost of player i on $\sigma'_i[u, w] \setminus p'$ in σ (recall that we consider the case where $\sigma'_i[u, w] \setminus p'$ does not contain any edges of σ_i). Since σ'_i is a best response of player i to σ_{-i} , the former cost is no less than the latter.

Using (2), we conclude that player j can improve her individual cost in σ by changing her path between u and w from p to $\sigma'_i[u, w]$, which contradicts the hy-

pothesis that σ_j is a best response of player j to σ . Formally,

$$\sum_{e \in p} d_e(\sigma_e) \geq \sum_{e \in p} d_e(\sigma'_e) > \sum_{e \in p'} d_e(\sigma'_e + 1) \geq \sum_{e \in \sigma'_i[u, w]} d_e(\sigma_e + 1)$$

The first inequality holds because $p \cap \sigma'_i = \emptyset$. Therefore, for all $e \in p$, $\sigma_e \geq \sigma'_e$, since the congestion of any edge not in σ'_i does not increase when player i switches from σ_i to σ'_i . The second inequality is (1), and the third inequality is (2).

If $\sigma'_i[u, w] \setminus p'$ contains some edges of σ_i , we show that due to the extension-parallel structure of the network, the congestion of the edges in p and p' does not change when player i switches from σ_i to σ'_i (see Fig. 1b). This contradicts the hypothesis that the best response of player j changes from σ_j to σ'_j when player i moves from σ_i to σ'_i .

We first show that the congestion of the edges of p does not change when player i switches from σ_i to σ'_i . Since u and w are connected by two internally disjoint paths p and $\sigma'_i[u, w]$, and since σ_i and $\sigma'_i[u, w]$ have some edges in common, by Proposition 1, σ_i contains u and w and does not have any edges in common with p . Hence the congestion of the edges of p is the same in both σ and σ' .

Next we show that the congestion of the edges of p' does not change when player i switches from σ_i to σ'_i . Let $\pi = \sigma'_i[u, w] \setminus p'$ and $\pi' = p' \setminus \sigma'_i[u, w]$ be the segments where $\sigma'_i[u, w]$ and p' deviate from each other. Since $\pi \cap \sigma_i \neq \emptyset$, π and π' are non-empty. Thus they are internally disjoint path segments with common endpoints, denoted u' and w' . Their first endpoint u' appears no sooner than u and their last endpoint w' appears no later than w on σ'_i and p' . Since σ_i has some edges in common with π , by Proposition 1, σ_i contains u' and w' and does not intersect π' at any vertex other than u' and w' . Moreover, since all three paths p' , $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ contain u' and w' , and since $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ deviate from p' between u' and w' , by Proposition 2, the three paths share their $u - u'$ and $w' - w$ segments (see also Fig. 1b). Therefore, if $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ deviate from each other, this happens between u' and w' . Consequently, the corresponding path segments $\sigma'_i[u', w']$ and $\sigma_i[u', w']$ do not contain any edges of p' . Thus the congestion of the edges of p' does not change when player i switches from σ_i to σ'_i .

Case Ib: $p \cap \sigma'_i \neq \emptyset$ We proceed with the case where p and σ'_i are not edge-disjoint. Then, by Proposition 1, σ'_i does not have any edges in common with p' and does not intersect p' at any vertex other than u and w . We first consider the case where $\sigma'_i[u, w] \setminus p$ does not contain any edges of σ_i (see e.g. Fig. 1c). Then,

$$\begin{aligned} & \sum_{e \in p \cap \sigma'_i} d_e(\sigma'_e) + \sum_{e \in p \setminus \sigma'_i} d_e(\sigma_e) \\ & \geq \sum_{e \in p} d_e(\sigma'_e) \\ & > \sum_{e \in p'} d_e(\sigma'_e + 1) \end{aligned}$$

$$\begin{aligned}
 &\geq \sum_{e \in p' \cap \sigma_i} d_e(\sigma_e) + \sum_{e \in p' \setminus \sigma_i} d_e(\sigma_e + 1) \\
 &\geq \sum_{e \in \sigma'_i[u, w]} d_e(\sigma'_e) \\
 &= \sum_{e \in p \cap \sigma'_i} d_e(\sigma'_e) + \sum_{e \in \sigma'_i[u, w] \setminus p} d_e(\sigma_e + 1) \tag{3}
 \end{aligned}$$

The first inequality holds because the congestion of any edge not in σ'_i does not increase when player i switches from σ_i to σ'_i (i.e. $\forall e \notin \sigma'_i, \sigma_e \geq \sigma'_e$). The second inequality is (1). The third inequality holds because when player i switches from σ_i to σ'_i : (i) the congestion of any edge decreases by at most 1 (i.e. $\forall e, \sigma'_e \geq \sigma_e - 1$), and (ii) the congestion of any edge not in σ_i does not decrease (i.e. $\forall e \notin \sigma_i, \sigma'_e \geq \sigma_e$). For the fourth inequality, we observe that the left-hand side is equal to the individual cost of player i on p' in σ (i.e. when the configuration of the remaining players is σ_{-i}), and that the right-hand side is equal to the cost of player i on $\sigma'_i[u, w]$ in σ . Since σ'_i is a best response of player i to σ_{-i} , the former cost is not less than the latter. The equality holds because $\sigma'_i[u, w] \setminus p$ does not contain any edges of σ_i and thus the congestion of every edge in $\sigma'_i[u, w] \setminus p$ increases by 1 when player i switches from σ_i to σ'_i (i.e. $\forall e \in \sigma'_i[u, w] \setminus p, \sigma'_e = \sigma_e + 1$). Thus (3) implies that

$$\sum_{e \in p \setminus \sigma'_i} d_e(\sigma_e) > \sum_{e \in \sigma'_i[u, w] \setminus p} d_e(\sigma_e + 1)$$

Therefore, player j can improve her individual cost in σ by switching from $p \setminus \sigma'_i$ to $\sigma'_i[u, w] \setminus p$. This contradicts the hypothesis that σ_j is a best response of player j to σ .

The case where $\sigma'_i[u, w] \setminus p$ contains some edges of σ_i (see e.g. Fig. 1d) is symmetric to the subcase of Ia where $\sigma'_i[u, w] \setminus p'$ contains some edge of σ_i . Applying the same arguments to the paths $\sigma'_i[u, w], \sigma_i[u, w]$ and p (instead of p' in the subcase of Ia), we show that the congestion of the edges in p and p' does not change when player i switches from σ_i to σ'_i . This contradicts the hypothesis that the best response of player j changes from σ_j to σ'_j when player i moves from σ_i to σ'_i . For completeness, we include the details in the [Appendix](#), Sect. A.2.

Case II: Either $u \notin \sigma'_i$ or $w \notin \sigma'_i$. We proceed with the case where σ'_i does not contain at least one of u and w . Then, by Proposition 1, σ'_i does have any edges in common with both p and p' .

If σ_i too does not contain at least one of u and w , then σ_i does not have any edges in common with both p and p' . Therefore, the congestion of the edges in p and p' does not change when player i switches from σ_i to σ'_i . This contradicts the hypothesis that the best response of player j changes from σ_j to σ'_j when player i moves from σ_i to σ'_i .

Hence we restrict our attention to the case where σ_i contains both u and w . Let $\sigma'_i \setminus \sigma_i$ and $\sigma_i \setminus \sigma'_i$ be the segments where σ_i and σ'_i deviate from each other. Since σ'_i does not contain either u or w (or both) and σ_i contains both u and w , $\sigma'_i \setminus \sigma_i$ and $\sigma_i \setminus \sigma'_i$ are non-empty. By Proposition 2, $\sigma'_i \setminus \sigma_i$ and $\sigma_i \setminus \sigma'_i$ are (non-empty)

internally disjoint path segments with common endpoints, denoted u' and w' . Their first endpoint u' appears no later than u and their last endpoint w' appears no sooner than w on σ_i and σ'_i (see e.g. Fig. 1e). Furthermore, either u is different from u' or w is different from w' (or both). By Proposition 2 and since σ_i deviates from at least one of p and p' between u and w , there is a unique path $\sigma_i[u', u]$ between u' and u and a unique path $\sigma_i[w, w']$ between w and w' . Let $z = \sigma_i[u', u] \cup \sigma_i[w, w']$. We highlight that both $\sigma_i[u', u]$ and $\sigma_i[w, w']$ are included in strategies σ_j and σ'_j . In particular, $\sigma_j[u', w'] = z \cup p$. Using the previous observations, we obtain that:

$$\begin{aligned}
 \sum_{e \in \sigma_j[u', w']} d_e(\sigma_e) &\geq \sum_{e \in z} d_e(\sigma_e) + \sum_{e \in p} d_e(\sigma'_e) \\
 &> \sum_{e \in z} d_e(\sigma_e) + \sum_{e \in p'} d_e(\sigma'_e + 1) \\
 &\geq \sum_{e \in z} d_e(\sigma_e) + \sum_{e \in p' \cap \sigma_i} d_e(\sigma_e) + \sum_{e \in p' \setminus \sigma_i} d_e(\sigma_e + 1) \\
 &\geq \sum_{e \in \sigma'_i[u', w']} d_e(\sigma_e + 1) \tag{4}
 \end{aligned}$$

The first inequality holds because the edges of p are not included in σ'_i , and the congestion of any edge not in σ'_i does not increase when player i moves from σ_i to σ'_i (i.e. $\forall e \notin \sigma'_i, \sigma_e \geq \sigma'_e$). The second inequality follows from (1). The third inequality holds because when player i switches from σ_i to σ'_i : (i) the congestion of any edge decreases by at most 1 (i.e. $\forall e, \sigma'_e \geq \sigma_e - 1$), and (ii) the congestion of any edge not in σ_i does not decrease (i.e. $\forall e \notin \sigma_i, \sigma'_e \geq \sigma_e$). For the fourth inequality, we observe that the left-hand side is equal to the individual cost of player i on $\sigma_i[u', u] \cup p' \cup \sigma_i[w, w']$ in σ (i.e. when the configuration of the remaining players is σ_{-i}), and that the right-hand side is equal to the individual cost of player i on $\sigma'_i[u', w']$ in σ (recall that $\sigma'_i[u', w']$ and $\sigma_i[u', w']$ are edge disjoint). Since σ'_i is a best response of player i to σ_{-i} , the former cost is not less than the latter.

Thus (4) implies that player j can decrease her individual cost in σ by switching from $\sigma_j[u', w']$ to $\sigma'_j[u', w']$. This contradicts the hypothesis that σ_j is a best response of player j to σ . Since we have reached a contradiction in all different cases, this concludes the proof of the lemma. □

By Lemma 1, once a player finds herself on her best response strategy, she will not have an incentive to deviate as long as the remaining players switch to their best improvement strategies. Hence we obtain the main result of this section:

Theorem 1 *For any n -player symmetric congestion game on an extension-parallel network, every best improvement sequence reaches a pure Nash equilibrium in at most n steps.*

4 Bounding the Inefficiency of Pure Nash Equilibria

We proceed to establish an upper bound of $\rho(\mathcal{D})$ on the PoS of symmetric network congestion games and on the PoA of symmetric games on extension parallel networks. We start with a technical property of the configurations of symmetric network congestion games that minimize the potential function Φ (cf. Lemma 2). Based on this property and on the definition of $\beta(\mathcal{D})$, we prove that for symmetric network games, the total cost of any global minimum of Φ is at most $(1 - \beta(\mathcal{D}))^{-1} = \rho(\mathcal{D})$ times the optimal total cost (cf. Lemma 3). Since a global minimum of Φ is also a local minimum, and thus a pure Nash equilibrium, we obtain that the PoS of symmetric network congestion games with latency functions in class \mathcal{D} is at most $\rho(\mathcal{D})$ (cf. Theorem 2). Furthermore, we show that for symmetric games on extension-parallel networks, every pure Nash equilibria is a *global* minimum of the potential function Φ (cf. Lemma 4). Therefore, for symmetric congestion games on extension-parallel networks with latency functions in class \mathcal{D} , the PoA is at most $\rho(\mathcal{D})$ (cf. Theorem 3).

Lemma 2 *Let Γ be a symmetric network congestion game, and let $G(V, E)$ be the underlying $s - t$ network. An acyclic configuration f of Γ minimizes the potential function Φ if and only if for every configuration g of Γ ,*

$$\Delta(f, g) \equiv \sum_{e: f_e > g_e} (f_e - g_e)d_e(f_e) - \sum_{e: f_e < g_e} (g_e - f_e)d_e(f_e + 1) \leq 0 \tag{5}$$

Proof For the if-part, we observe that for all configurations f, g ,

$$\begin{aligned} \Phi(f) - \Phi(g) &= \sum_{e: f_e > g_e} \sum_{i=g_e+1}^{f_e} d_e(i) - \sum_{e: f_e < g_e} \sum_{i=f_e+1}^{g_e} d_e(i) \\ &\leq \sum_{e: f_e > g_e} (f_e - g_e)d_e(f_e) - \sum_{e: f_e < g_e} (g_e - f_e)d_e(f_e + 1) = \Delta(f, g) \end{aligned}$$

Therefore, if for all configurations g of Γ , $\Delta(f, g) \leq 0$, then $\Phi(f) \leq \Phi(g)$ for all configurations g , and f is a global minimum of the potential function Φ .

For the converse, let f be an acyclic configuration that minimizes Φ , and let g be any acyclic configuration of Γ . The assumption that g is acyclic can be made without loss of generality, since any configuration g' can be turned into an acyclic configuration g with $g_e \leq g'_e$ for all $e \in E$, which implies that $\Delta(f, g') \leq \Delta(f, g)$. In the following, we assume that the configurations f and g are not identical, and consider the corresponding feasible $s - t$ flows.

Let $\hat{G}(V, \hat{E})$ be the network of the flow $g - f$. In particular, for each edge $(u, w) \in E$, \hat{E} contains a *forward* edge (u, w) with flow $g_{(u,w)} - f_{(u,w)}$ if $g_{(u,w)} > f_{(u,w)}$, a *backward* edge (w, u) with flow $f_{(u,w)} - g_{(u,w)}$ if $g_{(u,w)} < f_{(u,w)}$, and no edge between u and w if $f_{(u,w)} = g_{(u,w)}$. For a (directed) cycle C in \hat{G} , let

$$C^+ \equiv \{(u, w) \in E : (u, w) \in C \text{ and } g_{(u,w)} > f_{(u,w)}\}$$

be the set of edges of the original network G corresponding to the forward edges of C , let

$$C^- \equiv \{(u, w) \in E : (w, u) \in C \text{ and } g_{(u,w)} < f_{(u,w)}\}$$

be the set of edges of the original network G (with their directions as in E) corresponding to the backward edges of C , and let

$$\text{cost}_f(C) \equiv \sum_{e \in C^+} d_e(f_e + 1) - \sum_{e \in C^-} d_e(f_e)$$

be the cost (with respect to f) of cycle C . Next we establish the following two claims, which directly imply lemma: (i) if f is acyclic and a global minimum of the potential function Φ , then for every acyclic configuration g , the network \hat{G} of the flow $g - f$ does not contain any cycles C with $\text{cost}_f(C) < 0$, and (ii) if every cycle C in the network \hat{G} of the flow $g - f$ has $\text{cost}_f(C) \geq 0$, then $\Delta(f, g) \leq 0$.

We first establish claim (ii). Since f and g are acyclic $s - t$ flows of the same volume, a flow decomposition of $g - f$ yields only cycles and no paths of \hat{G} . Let $\{C_1, \dots, C_k\}$ be the set of (directed, simple) cycles of \hat{G} produced by the standard flow decomposition of $g - f$ (see e.g. the algorithm in [3, Theorem 3.5]), and let s_i be the amount of flow carried by each cycle C_i in the decomposition of $g - f$. Since g and f are acyclic $s - t$ flows, every cycle C_i contains at least one forward and at least one backward edge (i.e. both C_i^+ and C_i^- are non-empty).

By the properties of the standard flow decomposition, $\bigcup_{i \in [k]} C_i^+ = \{e \in E : g_e > f_e\}$ (i.e. the union of C_i^+ 's corresponds to the set of forward edges of \hat{G}), and $\bigcup_{i \in [k]} C_i^- = \{e \in E : g_e < f_e\}$ (i.e. the union of C_i^- 's corresponds to the set of backward edges of \hat{G}). Moreover, for every forward edge $(u, w) \in \hat{E}$, $\sum_{i: (u,w) \in C_i^+} s_i = g_{(u,w)} - f_{(u,w)}$, and for every backward edge $(w, u) \in \hat{E}$, $\sum_{i: (u,w) \in C_i^-} s_i = f_{(u,w)} - g_{(u,w)}$. Therefore,

$$\begin{aligned} \Delta(f, g) &= \sum_{e: f_e > g_e} (f_e - g_e)d_e(f_e) - \sum_{e: f_e < g_e} (g_e - f_e)d_e(f_e + 1) \\ &= \sum_{i=1}^k s_i \underbrace{\left(\sum_{e \in C_i^-} d_e(f_e) - \sum_{e \in C_i^+} d_e(f_e + 1) \right)}_{= -\text{cost}_f(C_i)} \\ &= - \sum_{i=1}^k s_i \text{cost}_f(C_i) \end{aligned} \tag{6}$$

Hence, if \hat{G} does not contain any (simple) cycles of negative cost (with respect to f), then $\Delta(f, g) \leq 0$. This concludes the proof of claim (ii).

We proceed to establish claim (i), namely that if f is a global minimum of the potential function Φ , then \hat{G} does not contain any cycles of negative cost (with respect to f). For sake of contradiction, we assume that there is a (directed) cycle C in \hat{G}

with $\text{cost}_f(C) < 0$. Given the flow f and the cycle C , we construct a $s - t$ flow f' of volume n by increasing the flow by one unit on the edges of C^+ , decreasing the flow by one unit on the edges of C^- , and keeping the same flow on the remaining edges. By construction, the potential value of f' is $\Phi(f') = \Phi(f) + \text{cost}_f(C) < \Phi(f)$, because $\text{cost}_f(C) < 0$. Since f and g are feasible acyclic $s - t$ flows, f' is a $s - t$ flow of volume n . In case where f' is not acyclic (and thus not feasible), we cancel the flow along all cycles in f' , and obtain a feasible acyclic $s - t$ flow f'' . Since this modification does not increase the amount of flow on any edge of the network, we have that $f''_e \leq f'_e$ for all $e \in E$, and that $\Phi(f'') \leq \Phi(f')$. Therefore, there exists a configuration of Γ of potential value less than $\Phi(f)$, which contradicts the hypothesis that f is a global minimum of Φ . This concludes the proof of claim (i), and the proof of the lemma. \square

The following lemma bounds the total cost of configurations (of symmetric network games) that minimize the potential function Φ in terms of the optimal total cost. The proof follows from the definition of $\rho(\mathcal{D})$ and Lemma 2.

Lemma 3 *Let Γ be a symmetric network congestion game with latency functions in class \mathcal{D} , let f be an acyclic configuration of Γ that minimizes the potential function Φ , and let o be the optimal configuration. Then, $C(f) \leq \rho(\mathcal{D})C(o)$.*

Proof Let $G(V, E)$ be the network underlying the definition of Γ . For every edge e with $f_e > o_e$,

$$\begin{aligned} f_e d_e(f_e) &= o_e d_e(f_e) + (f_e - o_e) d_e(f_e) \\ &\leq o_e d_e(o_e) + \beta(\mathcal{D}) f_e d_e(f_e) + (f_e - o_e) d_e(f_e) \end{aligned} \tag{7}$$

For the inequality, we use that $o_e(d_e(f_e) - d_e(o_e)) \leq \beta(\mathcal{D}) f_e d_e(f_e)$, which follows from the definition of $\beta(\mathcal{D})$, since $f_e > o_e$ and $d_e(x) \in \mathcal{D}$.

On the other hand, for every edge e with $f_e < o_e$,

$$\begin{aligned} f_e d_e(f_e) &= o_e d_e(o_e) - o_e d_e(o_e) + f_e d_e(f_e) \\ &\leq o_e d_e(o_e) - (o_e - f_e) d_e(f_e + 1) \end{aligned} \tag{8}$$

The inequality follows from $d_e(f_e) \leq d_e(f_e + 1)$ and $d_e(f_e + 1) \leq d_e(o_e)$, because the latency functions are non-decreasing and $f_e + 1 \leq o_e$ (recall that o_e and f_e are integral).

Using (7), (8), and Lemma 2, we obtain that:

$$\begin{aligned} C(f) &\leq C(o) + \beta(\mathcal{D}) \sum_{e: f_e > o_e} f_e d_e(f_e) \\ &= \overbrace{\Delta(f, o) \leq 0, \text{ by Lemma 2}} \\ &\quad + \sum_{e: f_e > o_e} (f_e - o_e) d_e(f_e) - \sum_{e: f_e < o_e} (o_e - f_e) d_e(f_e + 1) \\ &\leq C(o) + \beta(\mathcal{D}) C(f), \end{aligned}$$

which implies that $C(f) \leq (1 - \beta(\mathcal{D}))^{-1} C(o) = \rho(\mathcal{D})C(o)$. For the first inequality, we apply (7) to every edge e with $f_e > o_e$ and (8) to every edge e with $f_e < o_e$. For the last inequality, since f is an acyclic configuration that minimizes the potential function Φ and o is a configuration of Γ , Lemma 2 implies that $\Delta(f, o) \leq 0$. \square

Every symmetric network congestion game admits an acyclic configuration that minimizes the potential function Φ . Since the global minimum of Φ is also a local minimum, and thus a pure Nash equilibrium, Lemma 3 implies that every symmetric network game with latencies in class \mathcal{D} admits a pure Nash equilibrium of total cost at most $\rho(\mathcal{D})$ times the optimal total cost.

Theorem 2 *For any symmetric network congestion game with latency functions in class \mathcal{D} , the PoS is at most $\rho(\mathcal{D})$.*

For symmetric games on extension-parallel networks, configurations are acyclic and every pure Nash equilibrium is a *global* minimum of the potential function Φ (cf. Lemma 4). Therefore, by Lemma 3, for symmetric games on extension-parallel networks with latency functions in class \mathcal{D} , the total cost of any pure Nash equilibrium is at most $\rho(\mathcal{D})$ times the optimal total cost.

Theorem 3 *For symmetric congestion games on extension-parallel networks with latency functions in class \mathcal{D} , the PoA is at most $\rho(\mathcal{D})$.*

To conclude the proof of Theorem 3, we have to show that:

Lemma 4 *Every pure Nash equilibrium of a symmetric congestion game on an extension-parallel network is a global minimum of the potential function Φ .*

*Proof*⁷ Let Γ be a symmetric congestion game on an extension-parallel network $G(V, E)$, and let f be any pure Nash equilibrium of Γ . We show that for every configuration g of Γ , $\Delta(f, g) \leq 0$. Then by Lemma 2 (if-part), f is a global minimum of the potential function Φ . To establish that $\Delta(f, g) \leq 0$, we consider the network $\hat{G}(V, \hat{E})$ of the flow $g - f$, and show that if f is a pure Nash equilibrium, the cost (with respect to f) of any (directed) simple cycle in \hat{G} is non-negative (i.e. this corresponds to claim (i) in the proof of Lemma 2, only-if-part, claim (ii) holds with the same proof).

In the following, we use the same notation as in the proof of Lemma 2. We recall that $\Delta(f, g) = -\sum_{i=1}^k s_i \text{cost}_f(C_i)$ (see (6)), where $\{C_1, \dots, C_k\}$ is the set of (directed, simple) cycles produced by the standard flow decomposition of $g - f$, and s_i is the amount of flow carried by each cycle C_i in the decomposition of $f - g$ (recall that f, g are acyclic, as configurations of a symmetric game on an extension-parallel

⁷Lemma 4 is implicit in the work of Holzman and Law-Yone (see [19, Theorem 6.1]). However, the proof of Theorem 6.1 is omitted from [19]. Hence, for sake of completeness and because our framework is different from that of [19], we give a complete proof of Lemma 4 using techniques different from those in [19].

network). Therefore to establish that $\Delta(f, g) \leq 0$, it suffices to prove that for every (simple) cycle C_i in the decomposition of $g - f$, $\text{cost}_f(C_i) \geq 0$.

In fact, we prove that for every simple cycle C of \hat{G} , $\text{cost}_f(C) \geq 0$, which by the discussion above, implies the lemma. The crux of the proof is to show that due to extension-parallel structure of the network, every simple cycle C of \hat{G} contains two vertices u, w such that C^+ and C^- are two internally disjoint $u - w$ paths in the original network G (cf. Proposition 4). Then the claim follows from the hypothesis that f is a pure Nash equilibrium. More precisely, since C^- consists of backward edges only, $f_e > 0$ for all $e \in C^-$. Hence by Proposition 3, there is a player i whose strategy in f includes C^- . Therefore, $\sum_{e \in C^-} d_e(f_e) \leq \sum_{e \in C^+} d_e(f_e + 1)$, since otherwise player i could switch from C^- to C^+ between u and w and improve her individual cost, which contradicts the hypothesis that f is a pure Nash equilibrium. Consequently, for every simple cycle C of \hat{G} , $\text{cost}_f(C) \geq 0$.

To conclude the proof of the lemma, we have to show that the forward and the backward part of every cycle C in \hat{G} correspond to two internally paths in the extension-parallel network G .

Proposition 4 *Let G be any extension-parallel $s - t$ network, let f, g be any acyclic $s - t$ flows, let \hat{G} be the network of the flow $g - f$, and let C be any simple cycle of \hat{G} . Then there exist vertices u, w on C such that C^+ and C^- are two internally disjoint $u - w$ paths in G .*

Proof We first establish a property of \hat{G} required in the proof of Proposition 4.

Proposition 5 *Let G be any extension-parallel $s - t$ network, let f, g be any acyclic $s - t$ flows, and let \hat{G} be the network of the flow $g - f$. Then any (simple) $s - t$ path in \hat{G} consists of forward edges only, and any (simple) $t - s$ path in \hat{G} consists of backward edges only.*

Proof We prove only the first part, namely that any (simple) $s - t$ path in \hat{G} consists of forward edges only. The second part concerning $t - s$ paths follows by applying the first part to the flow $f - g$.

The proof is by induction on the extension-parallel structure of the network. The proposition holds if the network consists of a single edge (s, t) . We inductively assume that the proposition holds for any extension-parallel network with less edges than G .

If G is the parallel composition of two extension-parallel $s - t$ networks G_1 and G_2 , since g (resp. f) is an acyclic $s - t$ flow, it can be decomposed into an acyclic $s - t$ flow g^1 (resp. f^1) going through G_1 , and an acyclic $s - t$ flow g^2 (resp. f^2) going through G_2 . More precisely, g^1 (resp. f^1) is defined on G_1 only and is equal to g_e (resp. f_e) on every edge e of G_1 . Similarly, g^2 (resp. f^2) is defined on G_2 only and is equal to g_e (resp. f_e) on every edge e of G_2 . Let \hat{G}_1 (resp. \hat{G}_2) be the network of the flow $g^1 - f^1$ (resp. $g^2 - f^2$). Then \hat{G} is the parallel composition of \hat{G}_1 and \hat{G}_2 , and any $s - t$ path in \hat{G} goes through either \hat{G}_1 or \hat{G}_2 . The proposition follows from the induction hypothesis stating that $s - t$ paths in \hat{G}_1 and \hat{G}_2 consist of forward edges only.

If G is the series composition of an edge (s, s_1) and an extension-parallel $s_1 - t$ network G_1 , we assume that $g_{(s, s_1)} > f_{(s, s_1)}$ (i.e. that (s, s_1) is a forward edge in \hat{G}), since \hat{G} does not contain any $s - t$ path otherwise. Let g^1 (resp. f^1) be the flow obtained from g (resp. f) by ignoring the (flow on) edge (s, s_1) . Then g^1 (resp. f^1) is an acyclic $s_1 - t$ flow on G_1 with $g_e^1 = g_e$ (resp. $f_e^1 = f_e$) on every edge e of G_1 . Let \hat{G}_1 be the network of the flow $g^1 - f^1$. \hat{G} is the series composition of the forward edge (s, s_1) and \hat{G}_1 . Thus any $s - t$ path in \hat{G} begins with the forward edge (s, s_1) and continues with a $s_1 - t$ path in \hat{G}_1 . The proposition follows from the induction hypothesis stating that $s_1 - t$ paths in \hat{G}_1 consist of forward edges only. Exactly the same argument applies when G is the series composition of an extension-parallel $s - t_1$ network G_1 and an edge (t_1, t) . \square

We proceed to establish Proposition 4, namely that for every simple cycle C of \hat{G} , there are vertices u, w on C such that C^+ and C^- are two internally disjoint $u - w$ paths in G . The proof is by induction on the extension-parallel structure of the network. If the network consists of a single edge (s, t) , the proposition holds trivially because \hat{G} does not contain any cycles. We inductively assume that the proposition holds for any extension-parallel network with less edges than G .


If G is the series composition of an edge (s, s_1) and an extension-parallel $s_1 - t$ network G_1 , no cycle of \hat{G} contains s . Thus the vertex s and the edge (s, s_1) can be ignored. Let g^1 (resp. f^1) be the acyclic flow obtained from g (resp. f) by ignoring the (flow on) edge (s, s_1) , and let \hat{G}_1 be the network of the flow $g^1 - f^1$. Then every simple cycle C of \hat{G} is also a simple cycle of \hat{G}_1 (and vice versa), and the proposition follows from the induction hypothesis. The same argument applies when G is the series composition of an extension-parallel $s - t_1$ network G_1 and an edge (t_1, t) .

If G is the parallel composition of two extension-parallel $s - t$ networks G_1 and G_2 , since g (resp. f) is an acyclic $s - t$ flow, it can be decomposed into an acyclic $s - t$ flow g^1 (resp. f^1) going through G_1 , and an acyclic $s - t$ flow g^2 (resp. f^2) going through G_2 defined as in the proof of Proposition 5. Let \hat{G}_1 (resp. \hat{G}_2) be the network of the flow $g^1 - f^1$ (resp. $g^2 - f^2$). Then \hat{G} is the parallel composition of \hat{G}_1 and \hat{G}_2 . Therefore every simple cycle C of \hat{G} either lies entirely inside \hat{G}_1 (or \hat{G}_2) or consists of an $s - t$ path p^+ in \hat{G}_1 (resp. \hat{G}_2) followed by a $t - s$ path p^- in \hat{G}_2 (resp. \hat{G}_1). In the former case, the proposition follows from the induction hypothesis. In the latter case, by Proposition 5, the $s - t$ path p^+ consists of forward edges only and the $t - s$ path p^- consists of backward edges only. Hence $C^+ = p^+$ and $C^- = \{(u, w) \in E : (w, u) \in p^-\}$, and C^+ and C^- are two internally disjoint $s - t$ paths in G . \square

This concludes the proof of Lemma 4. \square

Remark The PoA may be greater than $\rho(\mathcal{D})$ even for series-parallel networks with linear latencies. For example, let us consider the 3-player game in Fig. 2. Since the latency functions are linear, $\rho = 4/3$. In the optimal configuration o , every edge has congestion 1 and the total cost is $C(o) = 11$ (e.g. an optimal configuration is obtained by assigning the first player to the upper path (e_1, e_3, e_5) , the second player

Fig. 2 A symmetric congestion game on a series-parallel network with linear latency functions and PoA greater than $4/3$



to the middle path (e_2, e_4, e_6), and the third player to the lower path (e_7). On the other hand, there is a pure Nash equilibrium f where the first player is assigned to (e_1, e_3, e_6) , the second player to (e_1, e_4, e_5) , and the third player to (e_2, e_3, e_5) . Each player incurs an individual cost of 5 and does not have an incentive to deviate to e_7 . The total cost is $C(f) = 15$ and the PoA is $15/11 > 4/3$. In addition, we observe that f is not a global minimum of the potential function Φ (this follows from Lemma 3 as well). In fact, $\Phi(f) = 12$, while the optimal configuration has $\Phi(o) = 11$ and is the global minimum of Φ (and thus a pure Nash equilibrium). Although Proposition 4 holds for the more general class of series-parallel networks, Lemma 4 fails in this example because Proposition 3 does not hold for series-parallel networks (and for the network in Fig. 2 in particular).

5 Conclusions

In this work, we considered symmetric congestion games on extension-parallel networks, an interesting and non-trivial generalization of parallel-link games, and established two remarkable properties previously known only for parallel-link games. In particular, we proved that for arbitrary (non-negative and non-decreasing) latency functions, any best improvement sequence reaches a pure Nash equilibrium in at most n steps, and that for latency functions in class \mathcal{D} , the PoA is at most $\rho(\mathcal{D})$. Moreover, our analysis implies that for symmetric network games with latency functions in class \mathcal{D} , the PoS is at most $\rho(\mathcal{D})$.

An interesting direction for further research is to investigate whether similar properties hold for more general classes of symmetric network congestion games. For instance, it would be quite interesting to establish a polynomial upper bound on the length of (certain) best improvement sequences for symmetric games on series-parallel networks. Lemma 1 may be a good starting point in this direction, since the analysis of several of the cases considered in the proof also applies to series-parallel networks. As for the inefficiency of Nash equilibria, it would be interesting to establish matching upper and lower bounds on the pure Price of Anarchy for symmetric congestion games on series-parallel networks and on general networks, especially if these bounds turn up to be considerably better than the known bounds for general congestion games (see e.g. [4, 6, 9]).

Appendix

A.1 The Proof of Proposition 1

We first recall that a $s - t$ network is series-parallel if and only if it does not contain a θ -network with terminals of degree 2 (see Fig. 3a) as a topological minor (see e.g. [24, Proposition 1].)

We observe that the proposition is trivial if $u = s$ and $w = t$. Hence, we assume that $u \neq s$. The case where $u = s$ and $w \neq t$ follows by symmetry.

Let π be a $s - t$ path sharing some edges with p . We first argue by contradiction that π does not intersect p' at any vertex other than u and w . Let us assume that π intersects p' at some vertex other than u and w . Since π has some edges in common with p , there is a segment of π connecting a vertex of $p - \{u, w\}$ to a vertex of $p' - \{u, w\}$ (or connecting a vertex of $p' - \{u, w\}$ to a vertex of $p - \{u, w\}$ depending on whether π first goes through an edge of p and then intersects $p' - \{u, w\}$ or the other way around). The endpoints of this segment of π belong to $(p \cup p') - \{u, w\}$. Thus the endpoints of this segment of π together with p and p' form a θ -network with terminals u and w , a contradiction. Consequently π does not intersect p' except possibly at u and w .

We then prove that π contains both u and w . To reach a contradiction, we assume that π does not contain u . Let π' be a path connecting s to u , and let $s', s' \neq u$, be the last common vertex of π and π' before u (s' exists because π and π' have a common source and π does not contain u). Let also $u', u' \neq u$, be the first common vertex of π and p after u (u' exists because π intersects p and π does not contain u). Then, the vertices s', u, u' , and w along with π' (the segment from s' to u), p (the segment from u to u' and the segment from u' to w), p' (the segment from u to w), and π (the segment from s' to u') form the θ -network of Fig. 3b, which is a topological minor of G . This contradicts the hypothesis that G is series-parallel network. If π contains u but not w , w is different from t . Then the proposition follows from the previous argument by symmetry (see also Fig. 3c).

A.2 Proof of Lemma 1: Case Ib, $\sigma'_i \cap p \neq \emptyset$ and $(\sigma'_i[u, w] \cap p) \cap \sigma_i \neq \emptyset$

We first show that the congestion of the edges of p' does not change when player i moves from σ_i to σ'_i . We recall that by Proposition 1, σ'_i does not have any edges in




Fig. 3 (a) The θ -network whose absence characterizes series-parallel $s - t$ networks. (b) and (c) If a path π has some edges in common with p (here p and π share the segment (u', w) in (b) and the segment (u, w') in (c)) and avoids either u or w , it forms (together with p , p' , and their extension to the source in (b) and to the sink in (c)) a network that includes the θ -network as a topological minor, and thus it is not series-parallel

common with p' and does not intersect p' at any vertex other than u and w . In addition, since u and w are connected by two internally disjoint paths p' and $\sigma'_i[u, w]$, and since σ_i and $\sigma'_i[u, w]$ have some edges in common, by Proposition 1, σ_i also contains u and w and does not have any edges in common with p' . Hence the congestion of the edges of p' is the same in both σ and σ' .

Next we prove that the congestion of the edges of p does not change when player i moves from σ_i to σ'_i . Let $\pi = \sigma'_i[u, w] \setminus p$ and $\pi' = p \setminus \sigma'_i[u, w]$ be the segments where $\sigma'_i[u, w]$ and p deviate from each other. Since $\pi \cap \sigma_i \neq \emptyset$, π and π' are non-empty. Thus π and π' are internally disjoint path segments with common endpoints, denoted u' and w' . Their first endpoint u' appears no sooner than u and their last endpoint w' appears no later than w on σ'_i and p . Since σ_i has some edges in common with π , by Proposition 1, σ_i contains u' and w' and does not intersect π' at any vertex other than u' and w' . Furthermore, since all three paths p , $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ contain u' and w' , and since $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ deviate from p between u' and w' , by Proposition 2, the three paths share their $u - u'$ and $w' - w$ segments (see also Fig. 1d). Therefore, if $\sigma'_i[u, w]$ and $\sigma_i[u, w]$ deviate from each other, this happens between u' and w' . Consequently, the corresponding path segments $\sigma'_i[u', w']$ and $\sigma_i[u', w']$ do not contain any edges of p . Thus the congestion of the edges of p does not change when player i switches from σ_i to σ'_i .

References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. In: Proc. of the 47th IEEE Symp. on Foundations of Computer Science (FOCS '06), pp. 613–622 (2006)
2. Ackermann, H., Röglin, H., Vöcking, B.: Pure Nash equilibria in player-specific and weighted congestion games. In: Proc. of the 2nd Workshop on Internet and Network Economics (WINE '06). LNCS, vol. 4286, pp. 50–61 (2006)
3. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network Flows: Theory, Algorithms, and Applications. Prentice Hall, New York (1993)
4. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. In: Proc. of the 23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS '06). LNCS, vol. 3884, pp. 218–229 (2006)
5. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM J. Comput. (2009, to appear). Preliminary version in Proc. of the 45th IEEE Symp. on Foundations of Computer Science (FOCS '04), pp. 295–304 (2004)
6. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05), pp. 57–66 (2005)
7. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. (2007, submitted). Preliminary version in Proc. of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06). LNCS, vol. 4051, pp. 311–322 (2006)
8. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA '07), pp. 169–178 (2007)
9. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proc. of the 37th ACM Symp. on Theory of Computing (STOC '05), pp. 67–73 (2005)
10. Christodoulou, G., Koutsoupias, E.: On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Proc. of the 13th European Symposium on Algorithms (ESA '05). LNCS, vol. 3669, pp. 59–70 (2005)
11. Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Selfish routing in capacitated networks. Math. Oper. Res. **29**(4), 961–976 (2004)

12. Epstein, A., Feldman, M., Mansour, Y.: Efficient graph topologies in network routing games. In: Joint Workshop on Economics of Networked Systems and Incentive-Based Computing (NetEcon+IBC '07) (2007)
13. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. In: Proc. of the 8th ACM Conference on Electronic Commerce (EC '07), pp. 84–92 (2007)
14. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to Nash equilibria in load balancing. *ACM Trans. Algorithms* **3**(3) (2007)
15. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure Nash equilibria. In: Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04), pp. 604–612 (2004)
16. Fotakis, D.: Stackelberg strategies for atomic congestion games. In: Proc. of the 15th European Symposium on Algorithms (ESA '07). LNCS, vol. 4698, pp. 299–310 (2007)
17. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: Nash equilibria in discrete routing games with convex latency functions. In: Proc. of the 31th International Colloquium on Automata, Languages and Programming (ICALP '04). LNCS, vol. 3142, pp. 645–657 (2004)
18. Gairing, M., Lücking, T., Monien, B., Tiemann, K.: Nash equilibria, the price of anarchy and the fully mixed Nash equilibrium conjecture. In: Proc. of the 32th International Colloquium on Automata, Languages and Programming (ICALP '05). LNCS, vol. 3580, pp. 51–65 (2005)
19. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. *Games Econ. Behav.* **21**, 85–101 (1997)
20. Holzman, R., Law-Yone (Lev-tov), N.: Network structure and strong equilibrium in route selection games. *Math. Soc. Sci.* **46**, 193–205 (2003)
21. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q.: Fast and compact: A simple class of congestion games. In: 20th National Conference on Artificial Intelligence (AAAI '05), pp. 489–494 (2005)
22. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS '99). LNCS, vol. 1563, pp. 404–413 (1999)
23. Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. In: Proc. of the 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS '04). LNCS, vol. 2996, pp. 547–558 (2004)
24. Milchtaich, I.: Network topology and the efficiency of equilibrium. *Games Econ. Behav.* **57**, 321–346 (2006)
25. Milchtaich, I.: The equilibrium existence problem in finite network congestion games. In: Proc. of the 2nd Workshop on Internet and Network Economics (WINE '06). LNCS, vol. 4286, pp. 87–98 (2006)
26. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2**, 65–67 (1973)
27. Roughgarden, T.: The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.* **67**(2), 341–364 (2003)
28. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM* **49**(2), 236–259 (2002)
29. Suri, S., Tóth, C.D., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* **47**(1), 79–96 (2007)