



Inefficiency of Pure Nash Equilibria in Series-Parallel Network Congestion Games

Bainian Hao^(✉)  and Carla Michini 

Department of Industrial and Systems Engineering, University of Wisconsin-Madison,
Madison, WI, USA
`{bhao8,michini}@wisc.edu`

Abstract. We study the inefficiency of pure Nash equilibria in symmetric unweighted network congestion games defined over series-parallel networks. We introduce a quantity $y(\mathcal{D})$ to upper bound the Price of Anarchy (PoA) for delay functions in class \mathcal{D} . When \mathcal{D} is the class of polynomial functions with highest degree p , our upper bound is $2^{p+1} - 1$, which is significantly smaller than the worst-case PoA for general networks. Thus, restricting to symmetric games over series-parallel networks can limit the inefficiency of pure Nash equilibria. We also construct a family of instances with polynomial delay functions that have a PoA in $\Omega(2^p/p)$ when the number of players goes to infinity. Compared with the subclass of extension-parallel networks, whose worst-case PoA is in $\Theta(p/\ln p)$, our results show that the worst-case PoA quickly degrades from sub-linear to exponential when relaxing the network topology. We also consider an alternative measure of the social cost of a strategy profile as the maximum players' cost. We introduce a parameter $z(\mathcal{D})$ and we show that the PoA is at most $y(\mathcal{D})z(\mathcal{D})$, which for polynomial delays of maximum degree p is at most $2^{2p+1} - 2^p$. Compared to the worst-case PoA in general symmetric congestion games, which is in $p^{\Theta(p)}$, our results shows a significant improvement in efficiency. We finally prove that our previous lower bound in $\Omega(2^p/p)$ is still valid for this measure of social cost. This is in stark contrast with the PoA in the subclass of extension-parallel networks, where each pure Nash equilibrium is a social optimum.

Keywords: Congestion games · Series-parallel networks · Price of anarchy

1 Introduction

In a non-cooperative game, rational players act selfishly to maximize their utility. The players influence each other's behaviour, since the quality of each player's strategy depends on the other players' actions. The notion of Nash equilibrium, where no player can improve her cost by unilaterally changing strategy, is the best-known solution concept for predicting a stable outcome of a game. However, since the players act selfishly and independently in a non-cooperative fashion, a Nash equilibrium might be far from minimizing the social cost. The inefficiency

of a Nash equilibrium can be measured by comparing its social cost against the minimum social cost that could be achieved. Precisely, the Price of Anarchy (PoA), introduced by Koutsoupias and Papadimitriou [21], is the largest ratio between the cost of a Nash equilibrium and the minimum social cost.

In this paper, we study network congestion games, where each player aims at selecting a shortest path from an origin to a destination, but the cost of each edge is non-decreasing with respect to the total number of players using it. These games are commonly used to model problems in large-scale networks such as routing in communication networks and traffic planning in road networks [19, 25] and represent a simple, yet powerful paradigm for selfish resource sharing.

We focus on the inefficiency of *pure* Nash equilibria. Unlike (mixed) Nash equilibria, where each player selects a probability distribution on her strategy set, in a *pure Nash equilibrium* (PNE) each player selects exactly one strategy from her strategy set. Pure Nash equilibria are not guaranteed to exist in general, but congestion games always admit one [26]. We consider two measures of social cost: the *total cost*, which is the sum of all players' costs, and the *maximum cost*, which is the maximum cost of a player in a strategy profile.

Several variants of network congestion games have been studied in the literature, which depend on the combination of a number of parameters. While some parameters seem to only marginally affect the PoA, the impact that graph structure has on the PoA is still not completely understood. Aland et al. [1] leave as an open direction the problem of characterizing “what structures provide immunity against a high PoA and what structures cause it”.

The approaches that have been proposed for general network congestion games [1–3, 7], later unified in the smoothness framework of Roughgarden [29, 30], cannot be used to derive stronger bounds that hold in the presence of special network structures. The two main graph structures for which stronger bounds on the PoA has been provided are parallel-links networks [5, 6, 15, 16, 23, 32] and extension parallel networks [12]. In this paper, we focus on the larger class of two-terminal series-parallel networks, and we provide upper and lower bounds on the worst-case PoA for (atomic, unweighted, symmetric) network congestion games. These networks can be recognized in linear-time [33] and are relevant in many applications, such as for problems on electric networks, scheduling and compiler optimization. Previous works have highlighted some strong properties of network congestion games defined over series-parallel networks, such as the existence of strong equilibria [20] and optimal tolls [14, 24].

First, we consider the total players' cost and arbitrary delay functions. Let \mathcal{D} be a class of nonnegative and non-decreasing functions. We introduce a new parameter $y(\mathcal{D})$ defined as

$$y(\mathcal{D}) = \sup_{d \in \mathcal{D}, x \in \mathbb{N}^+} \frac{(x+1)d(x+1) - xd(x)}{d(x)}, \quad (1)$$

which intuitively can be used to upper bound by what percentage the cost of an edge increases when one more player uses the edge. Note that $y(\mathcal{D}) \geq 1$ because $d(x) = (x+1)d(x) - xd(x) \leq (x+1)d(x+1) - xd(x)$. Our main result shows that the worst-case PoA in series-parallel networks is at most $y(\mathcal{D})$.

Theorem 1. *In a symmetric (unweighted) network congestion game on a series-parallel (s, t) -network with delays functions in class \mathcal{D} , the PoA w.r.t. the total players' cost is at most $y(\mathcal{D})$.*

The above result has interesting implications when \mathcal{D} is the class of polynomial functions with nonnegative coefficients and highest degree p . We show that in this case $y(\mathcal{D})$ is at most $2^{p+1} - 1$. Our result significantly improves over the worst-case PoA of unweighted congestion games, that is in $\Theta(p/\ln p)^{p+1}$ [1]. We point out that this worst-case PoA is also attained by unweighted network congestion games [1], however the construction used in [1] requires asymmetry. In the full version of this paper [17] we consider symmetric congestion games on general networks and we provide a family of instances violating the upper bound of Theorem 1. Moreover, we derive a lower bound on the worst-case PoA in symmetric network congestion games defined over series-parallel networks.

Theorem 2. *The worst-case PoA w.r.t. the total players' cost of a symmetric (unweighted) network congestion game on a series-parallel (s, t) -network, where the delay functions are polynomials with non-negative coefficients and highest degree p , is at least*

$$\frac{1}{1 + l^2 \sqrt[p]{r} - rl - \sqrt[p]{r} + r}, \quad (2)$$

where $r = \left(\frac{2}{2^{p+1}-1}\right)^{\frac{2^p}{2^p-1}}$ and $l = \frac{1}{2}r^{1-\frac{1}{2^p}}$.

We finally prove that our lower bound is in $\Omega\left(\frac{2^p}{p}\right)$, thus also in $\Omega(2^{cp})$ for each $c \in (0, 1)$, which almost asymptotically matches the upper bound of $2^{p+1} - 1$. Since the worst-case PoA in extension-parallel networks (a subclass of series-parallel networks) is in $\Theta(p/\ln p)$ [12, 13], our result shows that the PoA dramatically increases when relaxing the network topology from extension-parallel to series-parallel.

Next, we consider measuring the social cost of a strategy profile as the maximum players' cost. This variant of the social cost expresses the goal that a central authority might have to maximize fairness by minimizing the cost of the most disadvantaged player. We first consider arbitrary delay functions. To bound the PoA in this setting, introduce a new parameter $z(\mathcal{D})$ defined as

$$z(\mathcal{D}) = \sup_{d \in \mathcal{D}, x \in \mathbb{N}^+} \frac{d(x+1)}{d(x)}. \quad (3)$$

We first prove that the worst-case PoA in series-parallel networks is at most $y(\mathcal{D})z(\mathcal{D})$.

Theorem 3. *In a symmetric (unweighted) network congestion game on a series-parallel (s, t) -network with delays functions in class \mathcal{D} , the PoA w.r.t. the maximum players' cost is at most $z(\mathcal{D})y(\mathcal{D})$.*

When \mathcal{D} is the class of polynomial functions with nonnegative coefficients and maximum degree p we obtain that $z(\mathcal{D})$ is upper bounded by 2^p , thus the PoA is

at most $2^{2p+1} - 2^p$. Since the worst-case PoA for general symmetric congestion games and polynomial delays is in $p^{\Theta(p)}$ [7], our result shows a significant drop of the PoA in series-parallel networks.

Finally we show that the lower bound on the PoA w.r.t. the total players' cost also yields a valid lower bound when considering the maximum players' cost. We say that a class of networks \mathcal{N} is *closed under series compositions* if the series composition of two networks G^1 and G^2 in \mathcal{N} still belongs to \mathcal{N} .

Theorem 4. *Let \mathcal{N} be a class of networks closed under series compositions and let G be a network in \mathcal{N} . Then the worst-case PoA with respect to the maximum social cost of a symmetric (unweighted) network congestion game defined over G is at least the worst-case PoA with respect to the total social cost.*

For series-parallel networks and polynomial delays with nonnegative coefficients and maximum degree p Theorem 4 implies that the worst-case PoA is in $\Omega(2^p/p)$. This is in stark contrast with the result of [10], establishing that the PoA in extension-parallel networks is 1, i.e., any PNE is also a social optimum w.r.t. the maximum players' cost. Thus, relaxing the network topology from extension-parallel to series-parallel dramatically increases the inefficiency of pure Nash equilibria. The reason for this is that the key graph operations that we need to allow are the series compositions, which are forbidden for extension-parallel networks.

1.1 Further Related Work

Total Cost. There is a rich literature concerning the PoA in network congestion games where the social cost is measured based on the players' total cost. Many variants of network congestion games arise from considering different parameters and their combinations. As we shall see, the impact that graph structure has on the inefficiency of pure Nash equilibria varies significantly based on the combination of these parameters.

The first distinction is between atomic and non-atomic congestion games. In *non-atomic* congestion games, the number of players is infinite and each player controls an infinitesimal amount of flow. For these games, Roughgarden [27] proved that the PoA is independent of the network structure and equal to $\rho(\mathcal{D})$, where ρ depends on the class of delay functions \mathcal{D} [31].

For *atomic* games, where each player controls a non-negligible amount of flow, network structure affects the PoA differently, depending on whether all the players have the same effect on congestion. In *weighted* congestion games, where the effect of each player on congestion is proportional to the player's weight, the worst-case PoA is already achieved by very simple networks consisting of only parallel links [4] when \mathcal{D} is the class of polynomial functions with nonnegative coefficients and highest degree p . In contrast, in *unweighted* congestion games the effect of network structure seems significant. For asymmetric congestion games defined over general networks and in the case where \mathcal{D} is the class of polynomial functions with nonnegative coefficients, Christodoulou and Koutsoupias [7] showed that the PoA is in $p^{\Theta(p)}$ (see also [2, 3]). Aland et al. [1] later obtained

exact values for the worst-case PoA. These exact values admit a lower bound of $\lfloor \phi_p \rfloor^{p+1}$ and an upper bound of ϕ_p^{p+1} , where $\phi_p \in \Theta(p/\ln p)$ is the unique nonnegative real solution to $(x+1)^p = x^{p+1}$. For symmetric congestion games the PoA is again $p^{\Theta(p)}$ [2, 3, 7]. The worst case PoA drops significantly in the presence of special structure. Lücking et al. [22, 23] studied symmetric congestion games on parallel links and proved that the PoA is $4/3$ for linear functions. Later Fotakis [12] extended this result by proving an upper bound of $\rho(\mathcal{D})$ for the larger class of extension parallel networks with delays in class \mathcal{D} . Moreover, this upper bound is tight [11, 13]. It is known that, for the class of polynomial delays with nonnegative coefficients and highest degree p , $\rho(\mathcal{D}) \in \Theta(p/\ln p)$. This indicates that there is a huge gap between the worst-case PoA in general networks and in extension-parallel networks.

The PoA in symmetric series-parallel network congestion games has been recently investigated only for the specific case of affine delay functions [18], and it has been shown that the worst-case PoA is between $27/19$ and 2 [18], which is strictly worse than the PoA of $4/3$ in extension-parallel networks [12], and strictly better than the PoA of $5/2$ in general networks [8]. One key step to prove the upper bound in [18] consists in using the following inequality introduced in [12]

$$\frac{\text{cost}(f)}{\rho(\mathcal{D})} \leq \text{cost}(o) + \Delta(f, o), \quad (4)$$

where $\text{cost}(f)$ and $\text{cost}(o)$ denote the total cost of a PNE flow f and of a social optimum flow o , respectively, and $\Delta(f, o)$, is a quantity that depends on the difference $o - f$. For series-parallel networks with affine delays, Hao and Michini [18] prove that $\Delta(f, o) \leq 1/4 \text{cost}(f)$. This approach cannot be further extended to polynomial delays of maximum degree p , because we would obtain $\Delta(f, o) \leq \alpha(p) \text{cost}(f)$, where $\alpha(p)$ is a function of p that exceeds $1/\rho(\mathcal{D})$ for large p . Thus, an extension of the approach in [18] would provide an inconsequential bound.

Maximum Cost. The PoA with respect to the maximum players' cost has received less attention. In the non-atomic setting, Roughgarden [28] showed that the PoA is $n - 1$, where n is the number of nodes in the network.

In the atomic setting, Koutsoupias and Papadimitriou [21] first studied weighted congestion games with linear delay functions on m parallel links. For these games, they provided a lower bound of the PoA of $\Omega\left(\frac{\log m}{\log \log m}\right)$ and an upper bound of $O(\sqrt{m \log m})$. Later Czumaj and Vöcking [9] established a tight bound of $\Theta\left(\frac{\log m}{\log \log \log m}\right)$. Christodoulou and Koutsoupias [7] investigated general unweighted congestion games. In the symmetric case, they showed that the PoA is $5/2$ for affine delays and $p^{\Theta(p)}$ for polynomial delays of maximum degree p . In the asymmetric case, for games with N players, they proved that the PoA is in $\Theta(\sqrt{N})$ for affine delays and in $\Omega(N^{\frac{p}{p+1}})$ and $O(N)$ for polynomial delays of maximum degree p .

Epstein et al. [10] characterized efficient network topologies, i.e., graph topologies such that, for any class of non-decreasing delay functions, every PNE is also a social optimum. For unweighted symmetric network congestion games

they established that extension-parallel networks are efficient, implying that on these networks the PoA is 1. They also proved that this result is tight, i.e., it does not hold when further relaxing the network topology.

2 Preliminaries

Notation. Let $G = (V, E)$ be an (s, t) -network, i.e., a network with source s and sink t . Directed paths will be simply referred to as paths. A path from node u to node v is called a (u, v) -path. We will only consider *simple* paths, i.e., paths that do not traverse any node multiple times. Paths and cycles of G are regarded as sequences of edges, thus we may for example write $e \in p$ for a path p . An (s, t) -flow is an assignment of values to the edges of G such that, at each node u other than s and t , the sum of the values of the edges entering u equals the sum of the values of the edges leaving u . The value of the (s, t) -flow is the sum of the values of the edges entering t . We say a path p is *contained* in an (s, t) -flow f if for all $e \in p$, we have $f_e > 0$. For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$.

Network Congestion Games. Let $G = (V, E)$ be an (s, t) -network. We consider a network congestion game on G with N players. The strategy set X^i of player i is the set \mathcal{P} of (s, t) -paths in G . Since all the players have the same origin and destination, their strategy sets all coincide with \mathcal{P} and the game is called *symmetric*. A *state* of the game is a strategy profile $P = (p^1, \dots, p^N)$ where $p^i \in \mathcal{P}$ is the (s, t) -path chosen by player i , for $i \in [N]$. The set of states of the game is denoted by $X = X^1 \times \dots \times X^N$. Each state $P = (p^1, \dots, p^N) \in X$ induces an (s, t) -flow $f = f(P) = \chi^1 + \dots + \chi^N$ of value N , where χ^i is the incidence vector of p^i for all $i \in [N]$. We say that the (s, t) -paths p^1, \dots, p^N are a *decomposition* of the (s, t) -flow f if they induce flow f . Note that an (s, t) -flow f of value N can correspond to several states, since there might be multiple decompositions of f into N (s, t) -paths.

For each $e \in E$ we have a nondecreasing delay function $d_e : [N] \rightarrow \mathbb{R}_{\geq 0}$. Each player using e incurs a cost equal to $d_e(f_e)$, i.e., the cost of e depends on the total number of players that use e in f . Since d_e is a nondecreasing function, $d_e(j+1) \geq d_e(j)$ for $j \in [N-1]$, which models the effect of congestion. We denote the cost of a path p in G with respect to a flow f by $\text{cost}_f(p) = \sum_{e \in p} d_e(f_e)$. Thus, the cost incurred by player i in state P is $\text{cost}_f(p^i)$. We also define $\text{cost}_f^+(p) = \sum_{e \in p} d_e(f_e + 1)$. Finally, the cost of flow f in G is denoted by $\text{cost}(f) = \sum_{e \in E} f_e d_e(f_e)$. The *total cost* of a state P , denoted by $\text{tot}(P)$, is the sum of all players' costs. Clearly $\text{tot}(P)$ coincides with the cost of the flow $f(P)$:

$$\text{tot}(P) = \sum_{i \in [N]} \text{cost}_{f(P)}(p^i) = \text{cost}(f(P)).$$

We also define the *maximum cost* of P , denoted by $\max(P)$ as the maximum cost of a player in P :

$$\max(P) = \max_{i \in [N]} \text{cost}_{f(P)}(p^i).$$

Pure Nash Equilibria and Social Optima. A *pure Nash equilibrium* (PNE) is a state $(p^1, \dots, p^i, \dots, p^N)$ inducing an (s, t) -flow f such that, for each $i \in [N]$ we have

$$\text{cost}_f(p^i) \leq \text{cost}_{\tilde{f}}(\tilde{p}^i) \quad \forall (p^1, \dots, \tilde{p}^i, \dots, p^N) \in X \text{ inducing } (s, t) \text{-flow } \tilde{f}.$$

A PNE represents a stable outcome of the game, since no player $i \in [N]$ can improve her cost if she unilaterally changes strategy by selecting a different (s, t) -path \tilde{p}^i . With a slight abuse of terminology, we say that an (s, t) -flow f is a PNE if there exists a PNE $P = (p^1, \dots, p^N) \in X$ such that $f = f(P)$, i.e., f is the flow induced by P . On the other hand, we are also interested in a *social optimum*. We consider two definitions of social optimum, which depend on whether we measure the cost of a state P according to $\text{tot}(P)$ or $\text{max}(P)$. In the first case, a social optimum is a state that minimizes $\text{tot}(P) = \text{cost}(f(P))$ over all the states $P \in X$. With a slight abuse of terminology, we say that an (s, t) -flow o is a social optimum if o minimizes $\text{cost}(g)$ over all integral (s, t) -flows g of value N . In the second case a social optimum is a state that minimizes $\text{max}(P)$ over all the states $P \in X$. In other words, the social optimum is a state where the maximum player's cost is minimized.

Price of Anarchy. To measure the inefficiency of pure Nash equilibria, we use the definition of (pure) Price of Anarchy. The (pure) *Price of Anarchy* (PoA) is the maximum ratio between the cost of a PNE and the cost of a social optimum. In other words, to compute the PoA we consider the “worst” PNE, i.e., a PNE whose cost is as large as possible. For simplicity, from now on we will refer to the pure PoA as PoA.

We consider two definitions of PoA, which depend on whether we measure the cost of a state P according to $\text{tot}(P)$ or $\text{max}(P)$. In the first case, the PoA is the maximum ratio $\frac{\text{cost}(f)}{\text{cost}(o)}$ such that o is a social optimum flow and f is a PNE flow. In the second case, the PoA is the maximum ratio $\frac{\text{max}(P_f)}{\text{max}(P_o)}$ such that P_o is a social optimum state and P_f is a PNE.

Series-Parallel Networks. An (s, t) -network is series-parallel if it consists of either a single edge (s, t) or of two series-parallel networks composed either in series or in parallel. The *parallel composition* of two networks G_1 and G_2 is an (s, t) -network obtained from the union of G_1 and G_2 by identifying the source of G_1 and the source of G_2 into s , and by identifying the sink of G_1 and the sink of G_2 into t . The *series composition* of G_1 and G_2 , denoted by $G_1 \circ G_2$, is an (s, t) -network obtained from the union of G_1 and G_2 by letting s be the source of G_1 , t be the sink of G_2 , and by identifying the sink of G_1 with the source of G_2 . We remark that series-parallel networks are a superclass of parallel-link networks and extension-parallel networks, for which the PoA has been previously studied. An (s, t) -network is extension-parallel if it consists of a single edge (s, t) or of an extension-parallel network and a single edge composed either in series or in parallel.

3 Total Cost

3.1 Upper Bound on the PoA

In this section, we prove the upper bound on the PoA stated in Theorem 1. First, we need to introduce some necessary notation and properties of series-parallel networks. In the following, we denote by f and o a PNE and a social optimum, respectively, of the series-parallel network congestion game. We consider the graph $G(o-f)$ introduced in [12]. Precisely, the node set of $G(o-f)$ is V , and the edge set is $E(o-f) = \{(u, v) : (e = (u, v) \in E \text{ and } o_e - f_e > 0) \text{ or } (e = (v, u) \in E \text{ and } o_e - f_e < 0)\}$. $G(o-f)$ is a collection of simple cycles $\{C_1, \dots, C_h\}$ such that each C_i carries s_i units of flow. For each $i \in [h]$, define $C_i^+ = \{e = (u, v) \in E : (u, v) \in C_i, o_e > f_e\}$ and $C_i^- = \{e = (u, v) \in E : (v, u) \in C_i, o_e < f_e\}$.

Recall the parameter $y(\mathcal{D})$ we have defined in Sect. 1. In the next four lemmas, we will assume that there exists an index $i \in [h]$ such that C_i^+ is an (s, t) -path, and we will prove that the PoA is at most $y(\mathcal{D})$. Later, we will relax this assumption. Observe that, by definition, C_i^+ is contained in o . In the next lemma, we prove that the cost of C_i^+ with respect to o is at least the average players' cost in the PNE f , that is, $\text{cost}(f)/N$.

Lemma 1. *If C_i^+ , $i \in [h]$, is an (s, t) -path, then $\text{cost}_o(C_i^+) \geq \text{cost}(f)/N$.*

Proof. The cost of C_i^+ with respect to flow o satisfies:

$$\text{cost}_o(C_i^+) = \sum_{e \in C_i^+} d_e(o_e) \geq \sum_{e \in C_i^+} d_e(f_e + 1) \geq \frac{\text{cost}(f)}{N}.$$

The first inequality holds since for every $e \in C_i^+$, we have $o_e \geq f_e + 1$. Next we show that the second inequality holds. Denote by P^* the set of N (s, t) -paths in the PNE inducing f . Clearly $\max\{\text{cost}_f(\pi) : \pi \in P^*\} \geq \frac{\text{cost}(f)}{N}$. By contradiction, suppose that $\sum_{e \in C_i^+} d_e(f_e + 1) < \frac{\text{cost}(f)}{N}$. We would obtain that $\max\{\text{cost}_f(\pi) : \pi \in P^*\} > \text{cost}_f^+(C_i^+)$, thus one player would prefer to change her strategy into C_i^+ . This contradicts the fact that f is a PNE. \square

In the next lemma, we contemplate adding one unit of flow on an arbitrary (s, t) -path p contained in o , and we lower bound the corresponding increase of the total cost. This will be crucial to derive a lower bound on $\text{cost}_o(p)$ that will be used to relate $\text{cost}(f)$ and $\text{cost}(o)$.

Lemma 2. *Suppose that there exists an index $i \in [h]$, such that C_i^+ is an (s, t) -path. Then every (s, t) -path p contained in o satisfies*

$$\sum_{e \in p} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)) \geq \frac{\text{cost}(f)}{N}.$$

Proof. We will prove this by contradiction. Assume that there is an (s, t) -path p contained in o such that

$$\sum_{e \in p} (o_e + 1)d_e(o_e + 1) - \sum_{e \in p} o_e d_e(o_e) < \frac{\text{cost}(f)}{N}. \quad (5)$$

We define a new state o' obtained from o by deviating one unit of flow from C_i^+ to p . Let $S = C_i^+ \cap p$. First, the cost difference between o' and o is

$$\begin{aligned} \text{cost}(o') - \text{cost}(o) &= \sum_{e \in C_i^+ \setminus S} ((o_e - 1)d_e(o_e - 1) - o_e d_e(o_e)) \\ &\quad + \sum_{e \in p \setminus S} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)). \end{aligned}$$

Observe that, since the delay functions are non-decreasing, we have $d_e(o_e - 1) \leq d_e(o_e)$ for all $e \in C_i^+$, thus

$$\begin{aligned} \text{cost}(o') - \text{cost}(o) &\leq \sum_{e \in C_i^+ \setminus S} ((o_e - 1)d_e(o_e) - o_e d_e(o_e)) \\ &\quad + \sum_{e \in p \setminus S} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)) \\ &= - \sum_{e \in C_i^+ \setminus S} d_e(o_e) + \sum_{e \in p \setminus S} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)). \end{aligned}$$

Moreover, we have $d_e(o_e + 1) \geq d_e(o_e)$ for all $e \in S$, thus

$$\begin{aligned} 0 &\leq \sum_{e \in S} (o_e + 1)(d_e(o_e + 1) - d_e(o_e)) \\ &= - \sum_{e \in S} d_e(o_e) + \sum_{e \in S} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)). \end{aligned}$$

By summing up these two inequalities we get

$$\text{cost}(o') - \text{cost}(o) \leq - \sum_{e \in C_i^+} d_e(o_e) + \sum_{e \in p} ((o_e + 1)d_e(o_e + 1) - o_e d_e(o_e)).$$

By Lemma 1, since C_i^+ is an (s, t) -path, we have $\text{cost}_o(C_i^+) = \sum_{e \in C_i^+} d_e(o_e) \geq \frac{\text{cost}(f)}{N}$. Thus, by (5) we obtain $\text{cost}(o') - \text{cost}(o) < 0$, which contradicts the fact that o is a social optimum. \square

By using Lemma 2, we can derive a lower bound on $\text{cost}_o(p)$ similar to the lower bound on $\text{cost}_o(C_i^+)$ stated in Lemma 1, but with an extra factor of $y(\mathcal{D})$.

Lemma 3. *Suppose there exists an index $i \in [h]$ such that C_i^+ is an (s, t) -path, and let P be any decomposition of o . Then for every $p \in P$,*

$$y(\mathcal{D}) \text{cost}_o(p) \geq \frac{\text{cost}(f)}{N}.$$

Proof. Since P is a decomposition of o , for each $p \in P$ we have $o_e > 0$ for all $e \in p$. Then we have

$$y(\mathcal{D}) \text{cost}_o(p) = \sum_{e \in p} y(\mathcal{D}) d_e(o_e) \geq \sum_{e \in p} ((o_e + 1) d_e(o_e + 1) - o_e d_e(o_e)) \geq \frac{\text{cost}(f)}{N},$$

where the first inequality follows the definition of $y(\mathcal{D})$ stated in Eq. (1) and the second inequality follows from Lemma 2. \square

Finally, under the assumption that there exists a path C_i^+ from s to t , we are ready to prove that the PoA is at most $y(\mathcal{D})$.

Lemma 4. *If there exists an index $i \in [h]$ such that C_i^+ is an (s, t) -path, then $\text{cost}(f) \leq y(\mathcal{D}) \text{cost}(o)$.*

Proof. By Lemma 3 we know that given an arbitrary decomposition P of the social optimal flow o , for all $p \in P$, we have $y(\mathcal{D}) \text{cost}_o(p) \geq \frac{\text{cost}(f)}{N}$. Then we can conclude that:

$$y(\mathcal{D}) \text{cost}(o) = \sum_{p \in P} y(\mathcal{D}) \text{cost}_o(p) \geq |P| \frac{\text{cost}(f)}{N} = \text{cost}(f),$$

where the last equality follows from the fact that $|P| = N$. This implies that $\text{cost}(f) \leq y(\mathcal{D}) \text{cost}(o)$. \square

We now relax the assumption that there exists a path C_i^+ from s to t . In order to do this, we will exploit the structure of series-parallel graphs. If G is series-parallel, it is known that for each $i \in [h]$ C_i^+ and C_i^- are two internally disjoint paths in G from a node u_i to a node v_i [12]. For each $i \in [h]$, we identify the pair of nodes u_i, v_i and we define

$$\begin{aligned} V_i &= \{w \in V : \text{there is a } (u_i, v_i)\text{-path containing } w\}, \\ E_i &= \{e \in E : \text{there is a } (u_i, v_i)\text{-path containing } e\}, \end{aligned}$$

and we let $\mathcal{L} = \{E_1, \dots, E_h\}$.

Lemma 5. *If G is series-parallel, then $\mathcal{L} = \{E_1, \dots, E_h\}$ is a laminar family.*

Proof. We prove this lemma by showing that if $E_i \cap E_j \neq \emptyset$ for some i and j in $[h]$, then $E_i \subseteq E_j$ or $E_j \subseteq E_i$. We proceed by induction on $|E|$.

The base case as $|E| = 2$. If the two edges of G are composed in series, then there are no cycles. If they are composed in parallel, then there is only one cycle, i.e., $i = j$, and $E_i = E_j = E$. This implies that the lemma holds for the base case. Now we assume that when $|E| \leq t$, the lemma holds. When $|E| = t + 1$, since G is series-parallel, it can be decomposed either in series or in parallel.

Suppose that G can be decomposed in series into G_1 and G_2 . We first show that E_i and E_j are both contained either in the edge set of G_1 or in the edge set G_2 . In fact, E_i cannot have edges both in G_1 and in G_2 , otherwise C_i^+ and

C_i^- would not be internally disjoint paths. Thus E_i is contained either in the edge set of G_1 or in the edge set G_2 . Similarly, E_j is contained either in the edge set of G_1 or in the edge set G_2 . Moreover, E_i and E_j cannot belong to different components, otherwise we would have $E_i \cap E_j = \emptyset$. Thus, E_i and E_j both belong to the same component. Assume without loss of generality that this is G_1 . Since the number of edges of G_1 is at most t , by the inductive hypothesis we obtain that $E_i \subseteq E_j$ or $E_j \subseteq E_i$, thus the claim is proven in this case.

Now suppose that G can be decomposed in parallel into G_1 and G_2 . If E_i and E_j are both contained either in the edge set of G_1 or in the edge set of G_2 , then by induction the claim holds. If E_i is contained in the edge set of one component, say G_1 , and E_j is contained in the edge set of the other component G_2 , then $E_i \cap E_j = \emptyset$, a contradiction. Thus at least one among E_i and E_j has edges both in G_1 and in G_2 . Without loss of generality, suppose E_i does. We prove that C_i^+ and C_i^- are (internally disjoint) (s, t) -paths. By contradiction, suppose that C_i^+ and C_i^- are (s_i, t_i) -paths such that $s_i \neq s$ or $t_i \neq t$. Note that s_i and t_i are either both in G_1 or both in G_2 . Suppose w.l.o.g. they are both in G_1 . Then each (s_i, t_i) -path cannot contain any edge in G_2 . Because C_i^+ and C_i^- are (s, t) -paths, by the definition of E_i , we have $E_i = E$. Thus we conclude that $E_j \subseteq E_i$, which proves the claim in this case. \square

By Proposition 1 in [12], if w and w' are two nodes in V_i such that there exist two internally disjoint (w, w') -paths p_1 and p_2 , then every (s, t) -path having an edge in common with p_1 contains both w and w' and intersects p_2 only at w and w' . This implies that each (s, t) -path going through u_i also goes through v_i . As a consequence, for each $i \in [h]$ the sub-vectors of f and o that are indexed by the edges of E_i , denoted by $f(E_i)$ and $o(E_i)$, respectively, both define (u_i, v_i) -flows in the subgraph $G_i = (V_i, E_i)$. Define a network congestion game on G_i , where each edge $e \in E_i$ has the same delay d_e as in G , and the number of players N_i is equal to the value of flow $f(E_i)$.

Lemma 6. *If G is series-parallel and E_i is a maximal set in \mathcal{L} , then in the network congestion game defined on G_i , $f(E_i)$ and $o(E_i)$ are a PNE flow and a social optimum flow, respectively.*

Proof. Let N_i be the flow value of $f(E_i)$. First we show that $o(E_i)$ also has value N_i . Recall that $G(o - f)$ is a collection of cycles $\{C_1, \dots, C_h\}$ and each C_i carries s_i units of flow. By the definition of $G(o - f)$ we can change f into o as follows: for $j \in [h]$, decrease the flow on C_j^- by s_j and increase the flow on C_j^+ by s_j . By Lemma 5 \mathcal{L} is a laminar family, thus for each $j \in [h]$, the paths C_j^- and C_j^+ are either both in G_i or neither of them in G_i , i.e., either $E_j \subseteq E_i$, or $E_j \cap E_i = \emptyset$. Thus, each step does not change the flow value on G_i . We can conclude that when the procedure ends, the flow value $o(E_i)$ equals the flow value of $f(E_i) = N_i$.

Next, we show that $f(E_i)$ is a PNE flow on G_i . By contradiction, suppose that $f(E_i)$ is not a PNE flow on G_i . This implies that in each decomposition of $f(E_i)$ into N_i (u_i, v_i) -paths there is always one player who can decrease her cost

by deviating her strategy to another (u_i, v_i) -path in G_i . This implies that in each decomposition of f into $N(s, t)$ -paths there is always one player that can unilaterally deviate and decrease her cost. This contradicts to that f is a PNE flow.

Finally, we show that $o(E_i)$ is a social optimum on G_i . By contradiction, suppose that there is another flow $o'(E_i)$ in G_i of value N_i such that $\text{cost}(o'(E_i)) < \text{cost}(o(E_i))$. Then we can construct a flow o'' such that $o''_e = o_e$ for all $e \in E \setminus E_i$ and $o''_e = o'_e$ for all $e \in E_i$. Then $\text{cost}(o'') < \text{cost}(o)$, contradicting the fact that o is the social optimum. \square

We now consider the graphs G_i , $i \in [h]$, having node set V_i and edge set E_i .

Lemma 7. *If G is series-parallel and E_i is a maximal set in \mathcal{L} , then*

$$\text{cost}(f(E_i)) \leq y(\mathcal{D}) \text{cost}(o(E_i)).$$

Proof. According to Lemma 6, the congestion game with N_i players on the two terminal-series parallel graph G_i is such that $f(E_i)$ is a PNE and $o(E_i)$ is a social optimum. Note that u_i and v_i are, respectively, the source and the sink of G_i . Since C_i^+ is a (u_i, v_i) -path, by Lemma 4 we conclude that the lemma holds.

We are finally ready to prove Theorem 1, i.e., in a symmetric network congestion game defined over a series-parallel network with delay functions in class \mathcal{D} , the PoA is at most $y(\mathcal{D})$.

Proof of Theorem 1. Consider the PNE flow f , the social optimum flow o and the laminar family \mathcal{L} defined previously in this section. We will prove that, since G is series-parallel, then $\text{cost}(f) \leq y(\mathcal{D}) \text{cost}(o)$. Let E_{C_1}, \dots, E_{C_l} be the maximal sets in \mathcal{L} and denote by $E(\mathcal{L})$ their union. We rewrite $\text{cost}(f)$ as follows.

$$\text{cost}(f) = \sum_{e \notin E(\mathcal{L})} f_e d_e(f_e) + \sum_{e \in E(\mathcal{L})} f_e d_e(f_e).$$

Note that for each edge $e \notin E(\mathcal{L})$ we have $f_e = o_e$. Moreover, E_{C_1}, \dots, E_{C_l} are a partition of $E(\mathcal{L})$, since they are maximal members of \mathcal{L} that are pairwise disjoint. Thus we can rewrite the above expression as

$$\begin{aligned} \text{cost}(f) &= \sum_{e \notin E(\mathcal{L})} o_e d_e(o_e) + \sum_{i=1}^l \sum_{e \in E_{C_i}} f_e d_e(f_e) \\ &\leq y(\mathcal{D}) \sum_{e \notin E(\mathcal{L})} o_e d_e(o_e) + y(\mathcal{D}) \sum_{i=1}^l \sum_{e \in E_{C_i}} o_e d_e(o_e) = y(\mathcal{D}) \text{cost}(o), \end{aligned}$$

where the inequality follows from the fact that $y(\mathcal{D}) \geq 1$ and from Lemma 7. \square

Let Poly- p be the class of polynomial delay functions with maximum degree p , which are of the form $\sum_{j=0}^p a_j x^j$, with $a_j \geq 0$ for $j = 0, \dots, p$.

Lemma 8. *For the class of polynomial delay functions Poly- p it holds that $y(\text{Poly-}p) \leq 2^{p+1} - 1$.*

Proof. By using the definition of $y(\text{Poly-}p)$ in (1) we have that for any $x \in \mathbb{N}^+$

$$\begin{aligned} y(\text{Poly-}p) &= \sup_{a_0, \dots, a_p \in \mathbb{R}_{\geq 0}, x \in \mathbb{N}^+} \frac{(x+1) \sum_{j=0}^p a_j (x+1)^j - x \sum_{j=0}^p a_j x^j}{\sum_{j=0}^p a_j x^j} \\ &= \sup_{a_0, \dots, a_p \in \mathbb{R}_{\geq 0}, x \in \mathbb{N}^+} \frac{\sum_{j=0}^p (a_j ((x+1)^{j+1} - x^{j+1}))}{\sum_{j=0}^p a_j x^j}. \end{aligned} \quad (6)$$

We now exploit the fact that given two collections of nonnegative real numbers b_0, \dots, b_p and c_0, \dots, c_p , we have

$$\frac{\sum_{j=0}^p b_j}{\sum_{j=0}^p c_j} \leq \max_{j=0, \dots, p} \frac{b_j}{c_j}.$$

As a consequence, we can upper bound (6) by

$$\max_{j \in \{0, \dots, p\}, x \in \mathbb{N}^+} \frac{(x+1)^{j+1} - x^{j+1}}{x^j}. \quad (7)$$

We now upper bound the numerator of the above expression as follows:

$$(x+1)^{j+1} - x^{j+1} = \sum_{k=0}^{j+1} \binom{j+1}{k} x^{j+1-k} - x^{j+1} \leq \sum_{k=1}^{j+1} \binom{j+1}{k} x^j,$$

where the inequality follows from the fact that $j+1 \geq 1$ and $x \in \mathbb{N}^+$. From (7) we then obtain

$$y(\text{Poly-}p) \leq \max_{j \in \{0, \dots, p\}} \sum_{k=1}^{j+1} \binom{j+1}{k} = \max_{j \in \{0, \dots, p\}} \sum_{k=0}^{j+1} \binom{j+1}{k} - 1 = 2^{p+1} - 1.$$

□

By Theorem 1 and Lemma 8 we obtain that the PoA of series-parallel network congestion games with polynomial delay functions with highest degree is p is at most $2^{p+1} - 1$.

3.2 Lower Bound

In this section, we illustrate how to construct a family of instances that asymptotically achieve the lower bound on the PoA stated in Theorem 2. This construction is an extension to polynomial delays of the construction proposed in [18] for affine delays. Let $\{q_1, \dots, q_N\}$ be an ordered sequence of positive numbers such that $\sum_{i=1}^N q_i = 1$ and $q_{i+1} = \frac{1}{2^p} \sum_{j=1}^i \frac{q_j}{i}$ for $i \in [N-1]$. Let $m \in [N-1]$. We define a new sequence $\{s_1, \dots, s_N\}$ by averaging $\{q_1, \dots, q_m\}$. Precisely,

$s_1 = \dots = s_m = \frac{\sum_{i=1}^m q_i}{m}$ and $s_j = q_j$ for $j \geq m+1$. We construct a series-parallel (s, t) -network G with delays in Poly- p , and an (s, t) -flow f of value N recursively. Let G_m be a single (s, t) -edge with flow f_m of value m and delay equal to $\frac{s_1 x}{m}$. For every $i \in [m, N-1]$, we construct G_{i+1} and f_{i+1} using G_i and f_i as follows: we compose in parallel G_i and a new (s, t) -edge with flow value 1 and delay function $s_{i+1}x^p$ and call the new network \tilde{G}_i and the new (s, t) -flow \tilde{f}_i . Next, we compose in series $i+1$ copies of \tilde{G}_i with flow \tilde{f}_i to get G_{i+1} and f_{i+1} . We also divide the delay functions by $i+1$. Then we set $f = f_N$. Finally we compose G_N in parallel with m new (s, t) -edges e_1, \dots, e_m with delay function $\frac{1}{N}x^p$ to get G . By construction, G is a series-parallel network with polynomial delay functions having non-negative coefficients and maximum degree p .

To prove Theorem 2, we first show that f is a PNE. Then we define a new (s, t) -flow h that is obtained from f by deviating $k \in [m]$ units of flows from the most expensive (s, t) -paths in f to the k parallel (s, t) -edges in G with delay function $\frac{1}{N}x^p$. The parameters r and l in (2) are defined as $r = \frac{m}{N}$, $l = \frac{k}{m}$. The complete proof of Theorem 2 is given in the full version of this paper [17].

We now argue that the worst case PoA is in $\Omega(2^p/p)$. By substituting the expression of l in the denominator of (2), we obtain

$$1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r = 1 - \frac{1}{4}r^{2 - \frac{1}{2^p}} + r - r^{\frac{1}{2^p}}. \quad (8)$$

Since $r, l \in [0, 1]$, we can upper bound the above expression with

$$\begin{aligned} 1 + r - r^{\frac{1}{2^p}} &= 1 + \left(\frac{2}{2^{p+1} - 1} \right)^{\frac{2^p}{2^p - 1}} - \left(\frac{2}{2^{p+1} - 1} \right)^{\frac{1}{2^p - 1}} \\ &\leq 1 + \left(\frac{2}{2^{p+1} - 1} \right)^{\frac{2^p}{2^p - 1}} - \left(\frac{1}{2^p} \right)^{\frac{1}{2^p - 1}} \leq 1 - \left(1 - \frac{1}{2^p} \right) \left(\frac{1}{2^p - 1} \right)^{\frac{1}{2^p - 1}}. \end{aligned}$$

Finally, we have that $\lim_{p \rightarrow \infty} \frac{1 - (1 - \frac{1}{2^p}) \left(\frac{2}{2^{p+1} - 1} \right)^{\frac{2^p}{2^p - 1}}}{\frac{p}{2^p}} = 0$, proving that (8) is in $O(p/2^p)$, which implies that when N goes to infinity the PoA is at least in $\Omega(2^p/p)$.

4 Maximum Cost

In this section, we measure the social cost of a state P as the maximum players' cost in P , and we derive an upper bound and a lower bound on the PoA with respect to this notion of cost. Recall that given any state P , $\text{tot}(P)$ is the total cost of P and $\max(P)$ is the maximum cost of a player in P .

We first prove the upper bound on the PoA stated in Theorem 3.

Proof of Theorem 3. Let P_o be the social optimum with respect to the total cost, and let P_δ be the social optimum with respect to the maximum cost. Let $P_f = \{p_f^1, \dots, p_f^N\}$ be an arbitrary PNE. We will show that $\max(P_f) \leq z(\mathcal{D})y(\mathcal{D})\max(P_\delta)$.

Because P_f is a PNE and $\max(P_f)$ is the cost of a player, we have $\max(P_f) \leq \text{cost}_f^+(p_f^i)$ for any $i \in [N]$. Moreover, by (3), we have $\text{cost}_f^+(p_f^i) \leq z(\mathcal{D}) \text{cost}_f(p_f^i)$. In other words, the most expensive path in P_f has cost no greater than $z(\mathcal{D})$ times the cost of any other path in P_f . Thus we can conclude that

$$N \cdot \max(P_f) \leq \sum_{i=1}^N \text{cost}_f^+(p_f^i) \leq z(\mathcal{D}) \sum_{i=1}^N \text{cost}_f(p_f^i) = z(\mathcal{D}) \text{tot}(f),$$

i.e., the most expensive path in P_f has cost no greater than $z(\mathcal{D})$ times the average players' cost in P_f . Moreover,

$$z(\mathcal{D}) \text{tot}(P_f) \leq z(\mathcal{D})y(\mathcal{D}) \text{tot}(P_o) \quad (9)$$

$$\leq z(\mathcal{D})y(\mathcal{D}) \text{tot}(P_{\hat{o}}) \quad (10)$$

$$\leq z(\mathcal{D})y(\mathcal{D})(N \cdot \max(P_{\hat{o}})). \quad (11)$$

Inequality (9) directly follows Theorem 1. Inequality (10) holds since P_o is the social optimum state with respect to the total cost, which implies that $\text{tot}(P_o) \leq \text{tot}(P_{\hat{o}})$. Inequality (11) holds because $\max(P_{\hat{o}})$ is the maximum player's cost in $P_{\hat{o}}$. \square

We now consider the class Poly- p of polynomial delays with nonnegative coefficients and maximum degree p , and we prove that $z(\text{Poly-}p)$ is at most 2^p .

Lemma 9. *For the class of polynomial delay functions Poly- p it holds that $z(\text{Poly-}p) \leq 2^p$.*

Proof. By the definition of $z(\text{Poly-}p)$ in (3) we have that for any $x \in \mathbb{N}^+$

$$z(\text{Poly-}p) = \max_{x \in \mathbb{N}^+} \frac{\sum_{j=0}^p a_j (x+1)^j}{\sum_{j=0}^p a_j x^j}$$

Note that given two collections of nonnegative real numbers b_0, \dots, b_p and c_0, \dots, c_p , we have

$$\frac{\sum_{j=0}^p b_j}{\sum_{j=0}^p c_j} \leq \max_{j=0, \dots, p} \frac{b_j}{c_j}.$$

Thus,

$$z(\text{Poly-}p) = \max_{x \in \mathbb{N}^+} \frac{\sum_{j=0}^p a_j (x+1)^j}{\sum_{j=0}^p a_j x^j} \leq \max_{x \in \mathbb{N}^+} \max_{j=0, \dots, p} \frac{a_j (x+1)^j}{a_j x^j} \leq 2^p.$$

\square

Finally, we prove that, for any class of delay functions, and as long as the network's structure is preserved under series compositions, any lower bound on the PoA with respect to the total social cost is also valid when measuring the social cost in terms of the maximum players' cost.

Proof of Theorem 4. We start with an instance of an atomic, unweighted, symmetric network congestion game on a (s, t) -network G , where P_f is a PNE, P_o is a social optimum with respect to the total players' cost, and the PoA is $\text{cost}(P_f)/\text{cost}(P_o)$. Our goal is to construct a new instance on a network G' , and to define a PNE $P_{f'}$ and a social optimum $P_{o'}$ with respect to the maximum players' cost, such that

$$\frac{\max(P_{f'})}{\max(P_{o'})} = \frac{\text{cost}(P_f)}{\text{cost}(P_o)}.$$

We construct G' as follows. First, let G_1, \dots, G_N be N duplicates of G and let G' be the (s, t) -network obtained by composing in series G_1, \dots, G_N . We remark that any graph structure possessed by G is still valid for G' , by our assumption. Let $P_f = \{p_f^1, \dots, p_f^N\}$ and $P_o = \{p_o^1, \dots, p_o^N\}$. For each $i \in [N]$ let $P_{f_i} = \{p_{f_i}^1, \dots, p_{f_i}^N\}$ and $P_{o_i} = \{p_{o_i}^1, \dots, p_{o_i}^N\}$ be the corresponding duplicates of P_f and P_o in G_i , respectively. For each player $i \in [N]$ we define the strategy $p_{f'}^i$ of player i in $P_{f'}$ by having the player choose the path $p_{f_j}^{j(i)}$ in G_j , where $j(i) = (i + N - 1) \bmod N$. For example, the strategy of player 2 in $P_{f'}$ is obtained by composing in series the paths $p_{f_1}^2, p_{f_2}^3, \dots, p_{f_{N-1}}^N, p_{f_N}^1$. Analogously, we define the strategy $p_{o'}^i$ of player i in $P_{o'}$ by having the player choose the path $p_{o_j}^{j(i)}$ in G_j . It can be checked that $P_{f'} = \{p_{f'}^1, \dots, p_{f'}^N\}$ is a PNE for the new instance defined on G' (otherwise we would contradict that f is a PNE in the original instance). Similarly, it can be checked that $P_{o'} = \{p_{o'}^1, \dots, p_{o'}^N\}$ is the social optimum in G' with respect to the total cost (otherwise we would contradict that o is a social optimum in the original instance).

Observe that, since in our construction we are permuting the players' strategies, all the players have the same cost, both in $P_{f'}$ and in $P_{o'}$. Moreover this cost is equal to $\text{tot}(P_f)$ in $P_{f'}$ and to $\text{tot}(P_o)$ in $P_{o'}$. Thus, $\max(P_{f'}) = \text{tot}(P_f)$ and $\max(P_{o'}) = \text{tot}(P_o)$. Now let \hat{f} and \hat{o} be the worst PNE and the social optimum in the new instance. We conclude that

$$\frac{\text{tot}(P_f)}{\text{tot}(P_o)} = \frac{\max(P_{f'})}{\max(P_{o'})} \leq \frac{\max(P_{\hat{f}})}{\max(P_{\hat{o}})},$$

which implies the statement of this theorem. \square

5 Conclusion

Our contributions fill a gap in the literature on the PoA of atomic, unweighted, symmetric network congestion games, which tackles either general networks, or very simple network structures, such as parallel-link networks and extension-parallel networks.

In this paper we have focused on symmetric games. The worst-case PoA for unweighted congestion games over general networks [1] is achieved in the asymmetric case. On the other hand, Bhavalkar et al. proved that PoA of symmetric (unweighted) congestion games is as large as in asymmetric ones [4]. What

impact does symmetry have in the presence of network structure? Consider the class of polynomial delays Poly- p . If we relax the symmetry assumption, the upper bound of Theorem 1 does not hold. In fact, the PoA in asymmetric congestion games defined over parallel-link networks is as large as in asymmetric congestion games defined over general networks [16]. What happens if we instead stay in the realm of symmetric *network* congestion games, with no assumption on the network structure? In the full version of this paper [17], we provide a construction that violates the upper bound of Theorem 1, even if only by one.

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006). https://doi.org/10.1007/11672142_17
2. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, pp. 57–66. STOC 2005, Association for Computing Machinery, New York, NY, USA (2005)
3. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. *SIAM J. Comput.* **42**(1), 160–177 (2013)
4. Bhawalkar, K., Gairing, M., Roughgarden, T.: Weighted congestion games: the price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Econ. Comput.* **2**(4), 1–23 (2014)
5. Bilò, V., Vinci, C.: On the impact of singleton strategies in congestion games. In: Pruhs, K., Sohler, C. (eds.) 25th Annual European Symposium on Algorithms (ESA 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 87, pp. 17:1–17:14. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017)
6. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. *Algorithmica* **61**(3), 606–637 (2011)
7. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, pp. 67–73. STOC 2005, Association for Computing Machinery, New York, NY, USA (2005)
8. Correa, J., de Jong, J., de Keijzer, B., Uetz, M.: The inefficiency of Nash and subgame perfect equilibria for network routing. *Math. Oper. Res.* **44**(4), 1286–1303 (2019)
9. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Trans. Algorithms* **3**(1), 1–7 (2007)
10. Epstein, A., Feldman, M., Mansour, Y.: Efficient graph topologies in network routing games. *Games Econ. Behav.* **66**(1), 115–125 (2009)
11. Fotakis, D.: Stackelberg strategies for atomic congestion games. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 299–310. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75520-3_28
12. Fotakis, D.: Congestion games with linearly independent paths: convergence time and price of anarchy. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 33–45. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-79309-0_5

13. Fotakis, D.: Stackelberg strategies for atomic congestion games. *Theor. Comp. Sys.* **47**(1), 218–249 (2010)
14. Fotakis, D., Spirakis, G.P.: Cost-balancing tolls for atomic network congestion games. *Internet Math.* **5**(4), 343–363 (2008)
15. Gairing, M., Lüking, T., Mavronicolas, M., Monien, B., Rode, M.: Nash equilibria in discrete routing games with convex latency functions. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) *ICALP 2004*. LNCS, vol. 3142, pp. 645–657. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-27836-8_55
16. Gairing, M., Schoppmann, F.: Total latency in singleton congestion games. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 381–387. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-77105-0_42
17. Hao, B., Michini, C.: Inefficiency of pure Nash equilibria in series-parallel network congestion games, November 2021. *Optimization Online*
18. Hao, B., Michini, C.: The price of anarchy in series-parallel network congestion games. *Math. Program.* 1–31 (2022). <https://doi.org/10.1007/s10107-022-01803-w>
19. Harker, P.T.: Multiple equilibrium behaviors on networks. *Transp. Sci.* **22**(1), 39–46 (1988)
20. Holzman, R., Monderer, D.: Strong equilibrium in network congestion games: increasing versus decreasing costs. *Int. J. Game Theory* **44**(3), 647–666 (2015)
21. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999*. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49116-3_38
22. Lüking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. In: Diekert, V., Habib, M. (eds.) *STACS 2004*. LNCS, vol. 2996, pp. 547–558. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24749-4_48
23. Lüking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. *Theoret. Comput. Sci.* **406**(3), 187–206 (2008)
24. Nickerl, J.: The minimum tollbooth problem in atomic network congestion games with unsplittable flows. *Theor. Comp. Sys.* **65**(7), 1094–1109 (2021)
25. Orda, A., Rom, R., Shimkin, N.: Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Netw.* **1**(5), 510–521 (1993)
26. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Internat. 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.: The maximum latency of selfish routing. In: *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 980–981. SODA 2004, Society for Industrial and Applied Mathematics, USA (2004)
29. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, pp. 513–522. STOC 2009, Association for Computing Machinery, New York, NY, USA (2009)
30. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *J. ACM* **62**(5), 1–42 (2015)
31. Roughgarden, T., Tardos, E.: How bad is selfish routing? *J. ACM* **49**(2), 236–259 (2002)
32. Suri, S., Tóth, C.D., Zhou, Y.: Selfish load balancing and atomic congestion games. In: *Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 188–195. SPAA 2004, Association for Computing Machinery, New York, NY, USA (2004)
33. Valdes, J., Tarjan, R., Lawler, E.: The recognition of series parallel digraphs. *SIAM J. Comput.* **11**(2), 298–316 (1982)