Inefficiency of pure Nash equilibria in network congestion games: the impact of symmetry and graph structure*

Bainian Hao
bhao8@wisc.edu

Carla Michini
carla.michini@wisc.edu

University of Wisconsin-Madison

Abstract

We study the inefficiency of pure Nash equilibria in symmetric unweighted network congestion games. We first explore the impact of symmetry on the worst-case PoA of network congestion games. For polynomial delay functions with highest degree p, we construct a family of symmetric congestion games over arbitrary networks which achieves the same worst-case PoA of asymmetric network congestion games. Next, we explore the impact of network structure within the class of symmetric network congestion games. For delay functions in class $\mathcal{D}$ we introduce a quantity $y(\mathcal{D})$ and we show that the PoA is at most $y(\mathcal{D})$ in games defined over series-parallel networks. Thus, for polynomial delays with highest degree p, the PoA cannot exceed $2^{p+1}-1$, which is significantly smaller than the worst-case PoA in games defined over arbitrary networks. Moreover, we prove that the worst-case PoA quickly degrades from sub-linear to exponential when relaxing the network topology form extension-parallel to series-parallel.

Finally, we consider measuring the social cost as the maximum players' cost. For polynomial delays of maximum degree p we show that the worst case PoA is significantly smaller than that of general symmetric congestion games, but dramatically larger than that of games defined over extension-parallel networks.

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.

*A previous version of this work appeared in the proceedings of the Eighteenth Conference on Web and Internet Economics (WINE 2022).

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 symmetry and graph structure have on the PoA is still not completely understood. For example, Bhawalkar et al. [4, 5] proved that symmetry does not impact the PoA in *general* congestion games, but leave as an open question to establish whether this is still true for *network* congestion games. Correa et al. [9] partially answer this question by showing that, for affine delay functions, the worst-case PoA of symmetric network congestion games is as large as that of asymmetric network congestion games, when the number of players goes to infinity. In this paper we establish that, for polynomial delay functions of highest degree p, the PoA of symmetric network congestion games is very close to the PoA in asymmetric network congestion games.

Moreover, 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 [8, 2, 3, 1], 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 [32, 16, 17, 23, 7, 6] and extension parallel networks [13]. On the other hand, as remarked in both [1] and [13], the PoA in the larger class of series-parallel networks is not well understood. In this paper, we focus on two-terminal series-parallel networks, and we provide upper and lower bounds on the worst-case PoA for (atomic, unweighted) symmetric network congestion games. Two-terminal series-parallel networks 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 [15, 24].

1.1 Our contributions

First, we consider the total players' cost. For the class of polynomial delay functions of highest degree p, Aland et al. proved that the worst case PoA in *asymmetric* network congestion game is in $[\lfloor\Phi_p\rfloor^{p+1}, \Phi_p^{p+1}]$, where $\Phi_p \in \Theta(p/\ln p)$ [1]. Our first goal is to determine whether, in *symmetric* network congestion games the same worst-case PoA can be achieved. We answer this question in the affirmative.

Theorem 1.1. *The PoA of symmetric network congestion games with polynomial delay functions of highest degree p is at least $\lfloor\Phi_p\rfloor^{p+1}$.*

Theorem 1.1 indicates that symmetry has little impact on the worst-case PoA of network congestion games.

Next, we focus on *symmetric* congestion games defined over *series-parallel* networks. Let $\mathcal{D}$ be a class of nonnegative and non-decreasing functions. We introduce a new parameter $y(\mathcal{D})$ defined as

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

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 1.2. *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 indicates a significant drop of the worst-case PoA, which decreases from $\Theta(p/\ln p)^{p+1}$ in symmetric games over arbitrary networks (by Theorem 1.1) to $O(2^{p+1})$ in symmetric games over series-parallel networks.

We also provide a lower bound on the worst-case PoA in symmetric congestion games defined over series-parallel networks.

Theorem 1.3. *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*

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

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 $\Theta(p/\ln p)$ [13, 14], 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

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

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

Theorem 1.4. *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 $p^{\Theta(p)}$ [8], 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 1.5. *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 1.5 implies that the worst-case PoA is in $\Omega(2^p/p)$. This is in stark contrast with the result of [11], 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.

1.2 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 [5] 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 [8] 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 $p^{\Theta(p)}$ [8, 2, 3]. 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 [13] 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 [12, 14]. 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 [13], and strictly better than the PoA of 5/2 in general networks [9]. One key step to prove the upper bound in [18] consists in using the following inequality introduced in [13]

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

where cost(f) and 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 [10] established a tight bound of $\Theta\left(\frac{\log m}{\log\log\log m}\right)$. Christodoulou and Koutsoupias [8] 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. [11] 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 $\{1,\ldots,n\}$.

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 $P = (p^1,\ldots,p^N)$ 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 $X = X^1 \times \cdots \times X^N$. Each state $P = (p^1,\ldots,p^N) \in X$ induces an (s,t)-flow $f = f(P) = \chi^1 + \cdots + \chi^N$ 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 $p^1,\ldots,p^N$ 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 $d_e : [N] \to \mathbb{R}_{\geq 0}$. 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 $(p^1, \ldots, p^i, \ldots, p^N)$ inducing an (s,t)-flow f such that, for each $i \in [N]$ we have

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

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 $P = (p^1, \ldots, p^N) \in X$ 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 $\max(P)$. 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 $\max(P)$ 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 $\max(P)$. 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 $\frac{\max(P_f)}{\max(P_o)}$ 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 Lower bound on the PoA in arbitrary networks

In this section, we will prove the lower bound on the PoA stated in Theorem 1.1. Specifically, we construct a family of symmetric network congestion game instances that asymptotically achieves this lower bound.

Let $g_p(x) = (x+1)^p - x^{p+1}$. We recall that Φ_p is the unique nonnegative real solution to $g_p(x) = 0$ [1]. Clearly, $g_p(0) = 1 > 0$. Moreover, by Lemma 5.2 in [1], we know that $g_p(x)$ has exactly one local maximum ξ in $\mathbb{R}_+$, is strictly increasing in $[0,\xi)$ and strictly decreasing in (ξ,∞). As a result, for $0 \le x \le \phi_p$ we have $g_p(x) \ge 0$ and for $x \ge \phi_p$ we have $g_p(x) \le 0$. Next, we prove two additional properties of ϕ_p.

Lemma 3.1. *For every positive integer p and for all $k = 0,1,\dots,p-1$, we have $(\Phi_{p-1}+1)^k \le \Phi_{p-1}^{k+1}$.*

Proof. We proceed by backward induction. The base case is $k = p-1$. In this case, by the definition of Φ_p we know that $(\Phi_{p-1}+1)^{p-1} = \Phi_{p-1}^p$. We now assume that the claim holds for all $k \in \{2,\dots,p-1\}$, and we prove that the claim also holds for $k-1$. We have

$$(\Phi_{p-1}+1)^k \le \Phi_{p-1}^{k+1}$$

By dividing both terms by $\Phi_{p-1} > 0$ we obtain:

$$(\Phi_{p-1}+1)^{k-1} = \frac{(\Phi_{p-1}+1)^k}{\Phi_{p-1}+1} \le \frac{(\Phi_{p-1}+1)^k}{\Phi_{p-1}} \le \frac{\Phi_{p-1}^{k+1}}{\Phi_{p-1}} = \Phi_{p-1}^k.$$

□

Lemma 3.2. *For every positive integer p we have $\Phi_{p-1} \le \Phi_p \le \Phi_{p-1}+1$.*

Proof. We first prove that $\Phi_{p-1} \le \Phi_p$ by showing that $g_p(\Phi_{p-1}) = (\Phi_{p-1}+1)^p - \Phi_{p-1}^{p+1} > 0$. Since $\Phi_{p-1} > 0$ for every positive integer p, we equivalently show

$$\frac{(\Phi_{p-1}+1)^p}{\Phi_{p-1}} - \frac{\Phi_{p-1}^{p+1}}{\Phi_{p-1}} > \frac{(\Phi_{p-1}+1)^p}{\Phi_{p-1}+1} - \frac{\Phi_{p-1}^{p+1}}{\Phi_{p-1}} = (\Phi_{p-1}+1)^{p-1} - \Phi_{p-1}^p = 0.$$

Next, we prove that $\Phi_p \le \Phi_{p-1}+1$. To this purpose, we show that $g_p(\Phi_{p-1}+1) < 0$ which implies $\Phi_{p-1}+1 > \Phi_p$. Precisely, our goal is to prove

$$(\Phi_{p-1}+2)^p < (\Phi_{p-1}+1)^{p+1}. \tag{5}$$

We rewrite the left-hand-side of (5) as follows.

$$\begin{aligned}(\Phi_{p-1}+1+1)^p &= \sum_{k=0}^{p}\binom{p}{k}(\Phi_{p-1}+1)^k \\ &= (\Phi_{p-1}+1)^p + \sum_{k=0}^{p-1}\binom{p}{k}(\Phi_{p-1}+1)^k \\ &\le (\Phi_{p-1}+1)^p + \sum_{k=0}^{p-1}\binom{p}{k}\Phi_{p-1}^{k+1} \\ &= (\Phi_{p-1}+1)^p + \sum_{k=0}^{p}\binom{p}{k}\Phi_{p-1}^{k+1} - \Phi_{p-1}^{p+1} \\ &= (\Phi_{p-1}+1)^p + \Phi_{p-1}(\Phi_{p-1}+1)^p - \Phi_{p-1}^{p+1} \\ &= (\Phi_{p-1}+1)^{p+1} - \Phi_{p-1}^{p+1} \\ &< (\Phi_{p-1}+1)^{p+1}\end{aligned}$$

where the first inequality holds by Lemma 3.1. □

Proof of Theorem 1.1. We provide an example of an unweighted symmetric congestion game with delays in Poly-p that is defined over a network that is *not* series-parallel, and whose PoA asymptotically goes to $\lfloor\Phi_p\rfloor^{p+1}$ when the number of players is large. In this construction there are $N \ge \max\{\lfloor\Phi_p\rfloor+1, 2\lfloor\Phi_p\rfloor-2\}$ players. Let p be a positive integer. The graph G has $N(\lfloor\Phi_p\rfloor+N)+2$ nodes: the source s, the sink t, and N rows of $\lfloor\Phi_p\rfloor+N$ nodes. The nodes in row i are denoted by $v_{i,0}, v_{i,1}, \dots, v_{i,\lfloor\Phi_p\rfloor+N-1}$. In the following, for two integers h and k we denote by $h+k$ their sum modulo N. The graph G has N arcs $a_i = (s, v_{i,0})$ and N arcs $b_i = (v_{i,\lfloor\Phi_p\rfloor+N-1}, t)$ for all $i \in [N]$, having delay 0.

For all $i \in [N]$ and $j \in [\lfloor\Phi_p\rfloor+N-1]$ there is an arc e_{ij} from $v_{i,j-1}$ to $v_{i,j}$. The delay function associated to the arc is x^p for $\lfloor\Phi_p\rfloor \le j \le N$ and $c_j x^p$ for $1 \le j \le \lfloor\Phi_p\rfloor-1$ and $N+1 \le j \le \lfloor\Phi_p\rfloor+N-1$. For each j such that $1 \le j \le \lfloor\Phi_p\rfloor-1$, the coefficient c_j is computed as the solution of

$$c_j(j+1)^p = \sum_{i=j}^{\lfloor\Phi_p\rfloor-1} c_i \cdot i^p + j \cdot \lfloor\Phi_p\rfloor^p,$$

and for each $N+1 \le j \le \lfloor\Phi_p\rfloor+N-1$ we set $c_j = c_{\lfloor\Phi_p\rfloor+N-j}$. Finally, for all $i \in [N]$, $j \in [\lfloor\Phi_p\rfloor, \lfloor\Phi_p\rfloor+N-2]$ there is an arc g_{ij} from $v_{i,j}$ to $v_{i+1,j-\lfloor\Phi_p\rfloor+1}$ of constant delay 0. Note that for all $h \in [N]$ the only edge going to $v_{h,0}$ is a_h and the only edge going out from $v_{h,N+\lfloor\Phi_p\rfloor-1}$ is b_h.

We define the state $P = \{P^1, \dots, P^N\}$ where the (s,t)-path P^i chosen by player i is

$$\begin{aligned}&a_i, e_{i,1}, e_{i,2}, \dots, e_{i,\lfloor\Phi_p\rfloor}, g_{i,\lfloor\Phi_p\rfloor}, \\ &e_{i+1,2}, e_{i+1,3}, \dots, e_{i+1,\lfloor\Phi_p\rfloor+1}, g_{i+1,\lfloor\Phi_p\rfloor+1}, \\ &\dots, \\ &e_{i+N-1,N}, e_{i+N-1,N+1}, \dots, e_{i+N-1,N+\lfloor\Phi_p\rfloor-1}, b_{i+N-1}\end{aligned}$$

which selects $\lfloor\Phi_p\rfloor$ consecutive edges in each row. Since each edge with delay function x^p is used by $\lfloor\Phi_p\rfloor$ players, and each edge with delay function $c_j x^p$ for $j \in [\lfloor\Phi_p\rfloor]-1$ is used by j players, we conclude that the total players' cost in P is equal to $N((N-\lfloor\Phi_p\rfloor+1)\cdot\lfloor\Phi_p\rfloor\cdot\lfloor\Phi_p\rfloor^p+2\cdot\sum_{j=1}^{\lfloor\Phi_p\rfloor-1} c_j\cdot j^p)$.

We also define a state $P^* = \{P^{*1}, \ldots, P^{*N}\}$ where player i selects the path P^{*i} that only traverses row i:

$$a_i, e_{i,1}, e_{i,2}, \ldots, e_{i,N+\lfloor \Phi_p \rfloor -1}, b_i.$$

Since each edge e_{ij} is used by only one player, the total players' cost in P^* is equal to $N((N - \lfloor \Phi_p \rfloor + 1) + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor -1} c_j)$. We remark that the cost of P^* is an upper bound on the cost of a social optimal state.

We will next prove that P is a PNE. This will imply that for $N \to \infty$ the PoA is at least $\lfloor \Phi_p \rfloor^{p+1}$. Specifically,

$$\begin{aligned}
\frac{\mathrm{cost}(P)}{\mathrm{cost}(SO)} \geq \frac{\mathrm{cost}(P)}{\mathrm{cost}(P^*)} & \\
&= \lim_{N \to +\infty} \frac{(N - \lfloor \Phi_p \rfloor + 1) \cdot \lfloor \Phi_p \rfloor \cdot \lfloor \Phi_p \rfloor^p + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor -1} c_j \cdot j^p}{(N - \lfloor \Phi_p \rfloor + 1) + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor -1} c_j} \\
&= \lim_{N \to +\infty} \frac{(N - \lfloor \Phi_p \rfloor + 1) \cdot \lfloor \Phi_p \rfloor \cdot \lfloor \Phi_p \rfloor^p}{(N - \lfloor \Phi_p \rfloor + 1) + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor -1} c_j} \\
&= \lfloor \Phi_p \rfloor^{p+1}
\end{aligned}$$

To show that P is a PNE, we prove that every player is not able to decrease her cost by deviating to another (s,t)-path. Because the players' strategies are symmetric, without loss of generality, we consider player 1, whose strategy is

$$P^1 = a_1, P^1_1, g_{1,\lfloor \Phi_p \rfloor}, P^1_2, g_{2,\lfloor \Phi_p \rfloor +1}, \ldots, P^1_i, g_{i,\lfloor \Phi_p \rfloor + i - 1}, \ldots, P^1_N, b_N.$$

where for $i \in [N]$

$$P^1_i = e_{i,i}, e_{i,i+1}, \ldots, e_{i,i+\lfloor \Phi_p \rfloor -1}$$

Let f denote the flow induced by strategy profile P. To show that P is a PNE, we need the following claim:

Claim 1. For $i \in [N]$, $j \in [N-1]$, we have $d_{e_{i,j}}(f_{e_{i,j}}+1) \geq \mathrm{cost}(P^1_j)$ and $d_{e_{i,\lfloor \Phi_p \rfloor + N - j}}(f_{e_{i,\lfloor \Phi_p \rfloor + N - j}} + 1) \geq \mathrm{cost}(P^1_{N-j+1})$.

Proof of claim. Let $i \in [N]$ and $j \in [N-1]$. We prove $d_{e_{i,j}}(f_{e_{i,j}} + 1) \geq \mathrm{cost}(P^1_j)$. To show $d_{e_{i,\lfloor \Phi_p \rfloor + N - j}}(f_{e_{i,\lfloor \Phi_p \rfloor + N - j}} + 1) \geq \mathrm{cost}(P^1_{N-j+1})$ the proof is analogous.
Case(i): $i \in [N]$ and $j \in [\lfloor \Phi_p \rfloor - 1]$
We have $d_{e_{i,j}}(x) = c_j \cdot x^p$ and $f_{e_{i,j}} = j$. Thus we obtain

$$\begin{aligned}
d_{e_{i,j}}(f_{e_{i,j}} + 1) &= c_j \cdot (j+1)^p \\
&= \sum_{i=j}^{\lfloor \Phi_p \rfloor -1} c_i \cdot i^p + j \cdot \lfloor \Phi_p \rfloor^p \\
&= \mathrm{cost}(P^1_j),
\end{aligned}$$

where the second equality follows the definition of c_j.
Case(ii): $i \in [N]$ and $j \in [\lfloor \Phi_p \rfloor, N - \lfloor \Phi_p \rfloor + 1]$

Figure 1: An instance from our construction with $N = 3$ and delay functions in Poly- p with $p = 2$ or 3.

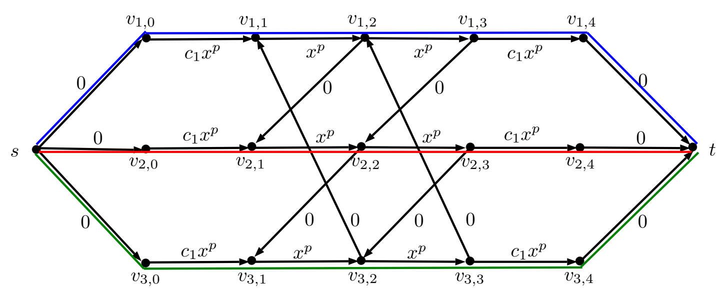

(a) Social optimum state

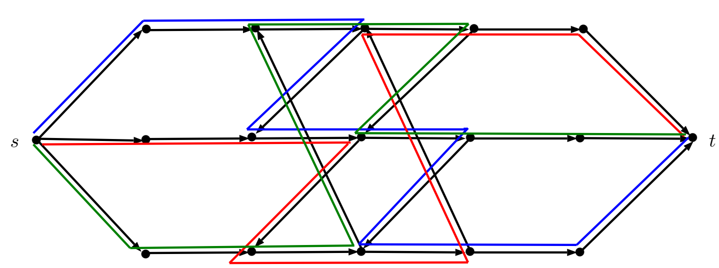

(b) PNE state

In this case, all the $\lfloor \Phi_p \rfloor$ edges in P_j^1 have cost x^p and are used by $\lfloor \Phi_p \rfloor$ players. As a consequence, $\mathrm{cost}(P_j^1) = \lfloor \Phi_p \rfloor^{p+1}$ and $d_{e_{i,j}}(f_{e_{i,j}} + 1) = (\lfloor \Phi_p \rfloor + 1)^p$. Since $0 \leq \lfloor \Phi_p \rfloor \leq \Phi_p$ we have $g_p(\lfloor \Phi_p \rfloor) \geq 0$ and we conclude that the claim holds.

Case(iii): $i \in [N]$ and $j \in [N - \lfloor \Phi_p \rfloor + 2, N - 1]$

Since $N \geq \max\{\lfloor \Phi_p \rfloor + 1, 2\lfloor \Phi_p \rfloor - 2\}$, we have $\lfloor \Phi_p \rfloor \leq j \leq N$, thus $e_{i,j}$ has delay x^p and is used by $\lfloor \Phi_p \rfloor$ players. We obtain

$$d_{e_{i,j}}(f_{e_{i,j}} + 1) = (\lfloor \Phi_p \rfloor + 1)^p \geq \lfloor \Phi_p \rfloor^{p+1},$$

where the inequality follows from the fact that $0 \leq \lfloor \Phi_p \rfloor \leq \Phi_p$ and $g_p(\lfloor \Phi_p \rfloor) \geq 0$. Next, we prove $\lfloor \Phi_p \rfloor^{p+1} \geq \mathrm{cost}(P_j^1)$. We recall that P_j^1 consists of $\lfloor \Phi_p \rfloor$ edges. In this case, the first $k = N - j + 1$ have delay x^p and are used by $\lfloor \Phi_p \rfloor$ players, while the last $\lfloor \Phi_p \rfloor - k$ edges are such that, for $\ell \in [\lfloor \Phi_p \rfloor - k]$ and $j = N + \ell$, edge (i, j) has delay $c_{N+\ell} x^p = c_{\lfloor \Phi_p \rfloor - \ell} x^p$ and is used by $\lfloor \Phi_p \rfloor - \ell$ players. Thus

$$\begin{aligned} \mathrm{cost}(P_j^1) &= \sum_{\ell=1}^{\lfloor \Phi_p \rfloor - k} c_{\lfloor \Phi_p \rfloor - \ell} \cdot (\lfloor \Phi_p \rfloor - \ell)^p + k \cdot \lfloor \Phi_p \rfloor^p \\ &= \sum_{\ell=k}^{\lfloor \Phi_p \rfloor - 1} c_\ell \cdot \ell^p + k \cdot \lfloor \Phi_p \rfloor^p. \end{aligned}$$

We next show that for $\ell \in [k, \lfloor \Phi_p \rfloor - 1]$ we have $c_\ell \cdot \ell^p \leq \lfloor \Phi_p \rfloor^p$, which directly implies our claim. We proceed by backward induction. First we show the base case, where $\ell = \lfloor \Phi_p \rfloor - 1$. By the definition of c_ℓ, we have

$$c_{\lfloor \Phi_p \rfloor - 1} \cdot \lfloor \Phi_p \rfloor^p = c_{\lfloor \Phi_p \rfloor - 1} \cdot (\lfloor \Phi_p \rfloor - 1)^p + (\lfloor \Phi_p \rfloor - 1) \cdot \lfloor \Phi_p \rfloor^p.$$

Thus we obtain

$$c_{\lfloor \Phi_p \rfloor - 1} = \frac{(\lfloor \Phi_p \rfloor - 1) \cdot \lfloor \Phi_p \rfloor^p}{\lfloor \Phi_p \rfloor^p - (\lfloor \Phi_p \rfloor - 1)^p}.$$

To prove our claim our goal is to show that

$$\frac{(\lfloor \Phi_p \rfloor - 1)^{p+1} \cdot \lfloor \Phi_p \rfloor^p}{\lfloor \Phi_p \rfloor^p - (\lfloor \Phi_p \rfloor - 1)^p} \leq \lfloor \Phi_p \rfloor^p,$$

or equivalently

$$\lfloor \Phi_p \rfloor^p \geq (\lfloor \Phi_p \rfloor - 1)^{p+1} + (\lfloor \Phi_p \rfloor - 1)^p = \lfloor \Phi_p \rfloor \cdot (\lfloor \Phi_p \rfloor - 1)^p.$$

We rewrite the above condition as

$$(\lfloor \Phi_p \rfloor - 1)^p \leq \lfloor \Phi_p \rfloor^{p-1}. \tag{6}$$

To prove that (6) is satisfied, we first observe that $\lfloor \Phi_{p-1} \rfloor^p \leq (\lfloor \Phi_{p-1} \rfloor + 1)^{p-1}$ because $0 \leq \lfloor \Phi_{p-1} \rfloor \leq \Phi_{p-1}$ implies $g_{p-1}(\lfloor \Phi_{p-1} \rfloor) \geq 0$. Moreover, $\Phi_p - \Phi_{p-1} \leq 1$ by Lemma 3.2, which implies $\lfloor \Phi_p \rfloor - \lfloor \Phi_{p-1} \rfloor \leq 1$ for all $p \in \mathbb{Z}^+$. Note that $\Phi_1 = \frac{1+\sqrt{5}}{2} > 1$ thus, by Lemma 3.2, $\Phi_p > 1$ for every positive integer p. Thus $0 \leq \lfloor \Phi_p \rfloor - 1 \leq \lfloor \Phi_{p-1} \rfloor \leq \Phi_{p-1}$ and $g_{p-1}(\lfloor \Phi_p \rfloor - 1) \geq 0$ implies (6).

Next, we assume our claim holds for $\ell \in [k+1, \lfloor \Phi_p \rfloor - 1]$ and we prove the claim still holds for $\ell - 1$. By definition we have

$$c_j (j+1)^p = \sum_{i=j}^{\lfloor \Phi_p \rfloor - 1} c_i \cdot i^p + j \cdot \lfloor \Phi_p \rfloor^p,$$

$$
\begin{aligned}
c_{\ell-1} \cdot \ell^p &= \sum_{h=\ell-1}^{\lfloor \Phi_p \rfloor -1} c_h \cdot h^p + (\ell-1) \cdot \lfloor \Phi_p \rfloor^p \\
&= c_{\ell-1} \cdot (\ell-1)^p + \sum_{h=\ell}^{\lfloor \Phi_p \rfloor -1} c_h \cdot h^p + (\ell-1) \cdot \lfloor \Phi_p \rfloor^p \\
&\leq c_{\ell-1} \cdot (\ell-1)^p + (\lfloor \Phi_p \rfloor - \ell) \lfloor \Phi_p \rfloor^p + (\ell-1) \cdot \lfloor \Phi_p \rfloor^p \\
&= c_{\ell-1} \cdot (\ell-1)^p + (\lfloor \Phi_p \rfloor - 1) \lfloor \Phi_p \rfloor^p .
\end{aligned}
$$

Thus

$$
c_{\ell-1} = \frac{(\lfloor \Phi_p \rfloor - 1) \cdot \lfloor \Phi_p \rfloor^p}{\ell^p - (\ell-1)^p}.
$$

To prove our claim our goal is to show that

$$
\frac{(\lfloor \Phi_p \rfloor - 1) \cdot \lfloor \Phi_p \rfloor^p \cdot (\ell-1)^p}{\ell^p - (\ell-1)^p} \leq \lfloor \Phi_p \rfloor^p ,
$$

that is equivalent to

$$
\lfloor \Phi_p \rfloor \cdot (\ell-1)^p \leq \ell^p,
$$

and in turn also to

$$
\left(\frac{\ell-1}{\lfloor \Phi_p \rfloor} \right)^p \cdot \lfloor \Phi_p \rfloor^{p+1} \leq \left(\frac{\ell}{\lfloor \Phi_p \rfloor + 1} \right)^p (\lfloor \Phi_p \rfloor + 1)^p. \tag{7}
$$

We conclude that (7) is satisfied since (i) $0 \leq \lfloor \Phi_p \rfloor \leq \Phi_p$, thus $g_p(\lfloor \Phi_p \rfloor) \geq 0$ implies $\lfloor \Phi_p \rfloor^{p+1} \leq (\lfloor \Phi_p \rfloor + 1)^p$ and (ii) $\frac{\ell-1}{\lfloor \Phi_p \rfloor} \leq \frac{\ell}{\lfloor \Phi_p \rfloor + 1}$ and since $\ell \leq \lfloor \Phi_p \rfloor + 1$. Thus Claim 1 is proved. $\diamond$

Now we show that player 1 is not able to strictly decrease their cost by deviating to another (s,t)-path. We provide an algorithm whose input is an arbitrary (s,t)-path $\bar{P}^1$ and output is another (s,t)-path with lower cost. We will show that by applying this algorithm repeatedly, the output will finally become P^1. First, we introduce some notations. We define the edge costs $w : E \to \mathbb{R}_{\geq 0}$ as $w_e = d_e(f_e)$ if $e \in P^1$ and $w_e = d_e(f_e + 1)$ if $e \notin P^1$. For any path q, let $w(q) = \sum_{e \in q} w_e$. Given a path q, let $q_{u,v}$, where $u, v \in q$, denotes the subpath of q between nodes u, v.

The algorithm receives in input an simple (s,t)-path $\bar{P}^1$ and returns in output another simple (s,t)-path $\hat{P}^1$ such that $w(\hat{P}^1) \leq w(\bar{P}^1)$. Let v_1 denote the last node in $\bar{P}^1$ such that $\bar{P}^1_{s,v_1}$ coincides with P^1_{s,v_1}. Note that we could have $v_1 = s$. If $v_1 = t$, then P^1 and $\bar{P}^1$ coincide. In this case the algorithm stops and returns P^1. If $v_1 \neq t$, then the algorithm determines the first node v_2 occurring after v_1 in $\bar{P}^1_{v_1,t}$, such that v_2 also belongs to P^1. Finally, the algorithm identifies the last node v_3 in $\bar{P}^1_{v_2,t}$ such that $\bar{P}^1_{v_2,v_3}$ coincides with $P^1_{v_2,v_3}$ and outputs $\hat{P}^1 = P^1_{s,v_3}, \bar{P}^1_{v_3,t}$. If $\hat{P}^1$ is not simple, then we make it simple by eliminating the cycles. If the algorithm returns a path $\hat{P}^1$ different from P^1, then the algorithm is applied again by setting as input $\bar{P}^1 = \hat{P}^1$.

We first argue that by repeatedly applying this algorithm we will finally obtain P^1 in output. In fact, either $v_1 = t$, or $\bar{P}^1_{v_3,t}$ is strictly contained in $\bar{P}^1_{v_1,t}$, since v_3 occurs after v_1 in $\bar{P}^1$. Thus, at the next iteration, when we set as input $\bar{P}^1 = \hat{P}^1$, the number of edges in the new $\bar{P}^1_{v_1,t}$ strictly decreases, since node v_1 of the current iteration coincides with node v_3 of the previous iteration. Next we show that $w(\hat{P}^1) \leq w(\bar{P}^1)$.

Case(i): $v_1 = s$. In this case,

$$
\bar{P}^1 = \bar{P}^1_{s,v_2}, P^1_{v_2,v_3}, \bar{P}^1_{v_3,t} \quad \text{and} \quad \hat{P}^1 = P^1_{s,v_2}, P^1_{v_2,v_3}, \bar{P}^1_{v_3,t}.
$$

Thus, we need to show that $w(\bar{P}^1_{s,v_2}) \geq w(P^1_{s,v_2})$. Since v_2 is the first node occurring after s in $\bar{P}^1$ that also belongs to P^1, by the definition of the weights w_e we have $w_e = d_e(f_e+1)$ for each $e \in \bar{P}^1_{s,v_2}$ and $w_e = d_e(f_e)$ for each $e \in P^1_{s,v_2}$. Suppose that $v_2 = v_{i,j}$. Then every path from s to $v_{i,j}$ must traverse an edge $e_{k,\ell}$ with $k \in [N]$ for each $\ell \in [j]$. Thus we have $w(\bar{P}^1_{s,v_{i,j}}) \geq \sum_{\ell=1}^{j} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1)$, where we arbitrarily picked $k = 1$ for each $\ell \in [j]$. By Claim 1

$$\sum_{\ell=1}^{\min\{N,j\}} \operatorname{cost}(P^1_\ell) \leq \sum_{\ell=1}^{j} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \leq w(\bar{P}^1_{s,v_{i,j}}). \tag{8}$$

Note that since $v_{i,j}$ belongs to P^1_i, it must be $i-1 \leq j \leq i + \lfloor \Phi_p \rfloor - 1$. If $j = i-1 \leq N$, then v_2 is the first node in P^1_i. Note that i cannot be 1 because the only edge goes into node $v_{1,0}$ is a_1, which belongs to P^1. If $v_2 = v_{1,0}$ then $a_1 \in P^1$ contradicts to the definition of v_2. Then we have

$$w(P^1_{s,v_{i,j}}) = \sum_{\ell=1}^{i-1} \operatorname{cost}(P^1_\ell) = \sum_{\ell=1}^{j} \operatorname{cost}(P^1_\ell). \tag{9}$$

If $i \leq j \leq i + \lfloor \Phi_p \rfloor - 1$, we have

$$w(P^1_{s,v_{i,j}}) \leq \sum_{\ell=1}^{i} \operatorname{cost}(P^1_\ell) \leq \sum_{\ell=1}^{\min\{N,j\}} \operatorname{cost}(P^1_\ell). \tag{10}$$

By combining equations (8), (9) and (10) we obtain that

$$w(P^1_{s,v_{i,j}}) \leq \sum_{\ell=1}^{\min\{N,j\}} \operatorname{cost}(P^1_\ell) \leq w(\bar{P}^1_{s,v_{i,j}}).$$

Case(ii): $v_2 = t$. In this case,

$$\bar{P}^1 = P^1_{s,v_1}, \bar{P}^1_{v_1,t} \quad \text{and} \quad \hat{P}^1 = P^1_{s,v_1}, P^1_{v_1,t} = P^1.$$

Thus, we need to show that $w(\bar{P}^1_{v_1,t}) \geq w(P^1_{v_1,t})$. Since $v_2 = t$ is the first node occurring after v_1 in $\bar{P}^1$ that also belongs to P^1, by the definition of the weights w_e we have $w_e = d_e(f_e+1)$ for each $e \in \bar{P}^1_{v_1,t}$ and $w_e = d_e(f_e)$ for each $e \in P^1_{v_1,t}$. Suppose that $v_1 = v_{i,j}$. Then every path from s to $v_{i,j}$ must traverse an edge $e_{k,\ell}$ with $k \in [N]$ for each $\ell \in [j+1, N + \lfloor \Phi_p \rfloor - 1]$. Thus we have $w(\bar{P}^1_{v_{i,j},t}) \geq \sum_{\ell=j+1}^{N+\lfloor \Phi_p \rfloor -1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1)$, where we arbitrarily picked $k = 1$ for each $\ell \in [j]$. By the second inequality of Claim 1

$$\sum_{\ell=\max\{j-\lfloor \Phi_p \rfloor+2,1\}}^{N} \operatorname{cost}(P^1_\ell) \leq \sum_{\ell=j+1}^{N+\lfloor \Phi_p \rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \leq w(\bar{P}^1_{s,v_{i,j}}). \tag{11}$$

Note that since $v_{i,j}$ belongs to P^1_i, it must be $i-1 \leq j \leq i + \lfloor \Phi_p \rfloor - 1$. If $j = i + \lfloor \Phi_p \rfloor - 1$, then v_1 is the last node of P^1_i. Note that i cannot be N because the only edge goes into node $v_{N,N+\lfloor \Phi_p \rfloor -1}$ is b_N, which belongs to P^1. If $v_1 = v_{N,N+\lfloor \Phi_p \rfloor-1}$ then $b_N \in P^1$ contradicts to the definition of v_1. Then we have

$$w(P^1_{v_{i,j},t}) = \sum_{\ell=i+1}^{N} \operatorname{cost}(P^1_\ell) \leq \sum_{\ell=\max\{j-\lfloor \Phi_p \rfloor+2,1\}}^{N} \operatorname{cost}(P^1_\ell). \tag{12}$$

If $i-1 \le j \le i+\lfloor \Phi_p \rfloor - 2$, we have

$$w(P^1_{v_{i,j},t}) \le \sum_{\ell=i}^{N} \text{cost}(P^1_\ell) \le \sum_{\ell=\max\{j-\lfloor \Phi_p \rfloor+2,1\}}^{N} \text{cost}(P^1_\ell). \tag{13}$$

By combining equations (11), (12) and (13) we obtain that

$$w(P^1_{v_{i,j},t}) \le \sum_{\ell=\max\{j-\lfloor \Phi_p \rfloor+2,1\}}^{N} \text{cost}(P^1_\ell) \le w(\bar{P}^1_{v_{i,j},t}).$$

Case(iii): $v_1 \neq s$ and $v_2 \neq t$. Without loss of generality, let $v_1 = v_{i,j}$ and $v_2 = v_{h,k}$. Note that v_1 belongs to P^1_i, thus $i-1 \le j \le i+\lfloor \Phi_p \rfloor - 1$, and v_2 belongs to P^1_h, thus $h-1 \le k \le h+\lfloor \Phi_p \rfloor - 1$. Finally, note that $k > j$ if $i = h$. Thus $\bar{P}^1 = (P^1_{s,v_1}, \bar{P}^1_{v_1,v_2}, P^1_{v_2,v_3}, \bar{P}^1_{v_3,t})$ and $\hat{P}^1 = (P^1_{s,v_1}, P^1_{v_1,v_2}, P^1_{v_2,v_3}, \bar{P}^1_{v_3,t})$. So we only need to show that

$$w(\bar{P}^1_{v_1,v_2}) \ge w(P^1_{v_1,v_2}).$$

Since v_2 is the first node occurring after s is $\bar{P}^1$ that also belongs to P^1, by the definition of the weights w_e we have $w_e = d_e(f_e+1)$ for each $e \in \bar{P}^1_{v_1,v_2}$ and $w_e = d_e(f_e)$ for each $e \in P^1_{v_1,v_2}$.

Subcase(iii).1: $j = i+\lfloor \Phi_p \rfloor - 1$, i.e., $v_{i,j}$ is the last node in P^1_i, and $h \le k \le h+\lfloor \Phi_p \rfloor - 1$. Thus $h > i$ and $h \neq k-1$, otherwise $\bar{P}^1_{v_1,v_2}$ will intersect with $P^1_{v_1,v_2}$ before v_2. Since $g_{i,j} \in P^1_{v_1,v_2}$, we conclude that the first edge in $\bar{P}^1_{v_1,v_2}$ is $e_{i,i+\lfloor \Phi_p \rfloor}$. And since $e_{h,k} \in P^1_{v_1,v_2}$, we conclude that the last edge in $\bar{P}^1_{v_1,v_2}$ is $g_{h-1,k+\lfloor \Phi_p \rfloor-1}$. Note that because $k+\lfloor \Phi_p \rfloor - 1 \ge h+\lfloor \Phi_p \rfloor - 1 \ge i+\lfloor \Phi_p \rfloor$, then every path begin with $e_{i,i+\lfloor \Phi_p \rfloor}$ and end with $g_{h-1,k+\lfloor \Phi_p \rfloor-1}$ must traverse an edge $e_{m,\ell}$ with $m \in [N]$ for each $\ell \in [i+\lfloor \Phi_p \rfloor, k+\lfloor \Phi_p \rfloor - 1]$. We arbitrarily picked $m = 1$ for each ℓ, then we can conclude that

$$w(\bar{P}^1_{v_1,v_2}) \ge \sum_{\ell=i+\lfloor \Phi_p \rfloor}^{k+\lfloor \Phi_p \rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \ge \sum_{\ell=i+\lfloor \Phi_p \rfloor}^{h+\lfloor \Phi_p \rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1). \tag{14}$$

Then according to the second inequality of Claim 1, we have

$$\sum_{\ell=i+\lfloor \Phi_p \rfloor}^{h+\lfloor \Phi_p \rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \ge \sum_{\ell=i+1}^{h} w(P^1_\ell). \tag{15}$$

Recall that $v_1 = v_{i,j}$ is the last node in P^1_i, which implies that $P^1_{v_1,v_2}$ contains subpath $P^1_{i+1}, \ldots, P^1_{h-1}$ and part of subpath P^1_h. So we can conclude that

$$\sum_{\ell=i+1}^{h} \text{cost}(P^1_\ell) \ge w(P^1_{v_1,v_2}) \tag{16}$$

By combining inequalities (14), (15) and (16) we have $w(\bar{P}^1_{v_1,v_2}) \ge w(P^1_{v_1,v_2})$.

Subcase(iii).2: $i-1 \le j \le i+\lfloor \Phi_p \rfloor - 2$ and $k = h-1$, i.e., $v_{h,k}$ is the first node in P^1_h. Thus we have $h > i$. Since $e_{i,j+1} \in P^1_{v_1,v_2}$, we conclude that the first edge in $\bar{P}^1_{v_1,v_2}$ is $g_{i,j}$. And since $g_{h-1,k+\lfloor \Phi_p \rfloor-1}$, who ends with $v_{h,k}$, belongs to $P^1_{v_1,v_2}$, we conclude that the last edge in $\bar{P}^1_{v_1,v_2}$ is $e_{h,h-1}$. Note that because $g_{i,j}$ ends with node $v_{i+1,j-\lfloor \Phi_p \rfloor+1}$ and $h-1 \ge j-\lfloor \Phi_p \rfloor + 2$, then every path begin with $g_{i,j}$

and end with $e_{h,h-1}$ must traverse an edge $e_{m,\ell}$ with $m \in [N]$ for each $\ell \in [j - \lfloor \Phi_p \rfloor + 2, h-1]$. We arbitrarily picked $m = 1$ for each ℓ, then we can conclude that

$$w(\bar{P}^1_{v_1,v_2}) \geq \sum_{\ell=j-\lfloor \Phi_p \rfloor+2}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1). \tag{17}$$

Then according to the first inequality of Claim 1, we have

$$\sum_{\ell=i}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h-1} w(P^1_\ell). \tag{18}$$

Recall that $v_2 = v_{h,k}$ is the first node in P^1_i, which implies that $P^1_{v_1,v_2}$ contains subpath $P^1_{i+1}, \dots, P^1_{h-1}$ and part of subpath P^1_i. So we can conclude that

$$\sum_{\ell=i}^{h-1} \text{cost}(P^1_\ell) \geq w(P^1_{v_1,v_2}) \tag{19}$$

By combining inequalities (17), (18) and (19) we have $w(\bar{P}^1_{v_1,v_2}) \geq w(P^1_{v_1,v_2})$.
Subcase(iii).3: $i - 1 \leq j \leq i + \lfloor \Phi_p \rfloor - 2$ and $h \leq k \leq h + \lfloor \Phi_p \rfloor - 1$. Thus we have $h \geq i$ and $k > j$ if $h = i$. Since $e_{i,j+1} \in P^1_{v_1,v_2}$, we conclude that the first edge in $\bar{P}^1_{v_1,v_2}$ is $g_{i,j}$. And since $e_{h,k} \in P^1_{v_1,v_2}$, we conclude that the last edge in $\bar{P}^1_{v_1,v_2}$ is $g_{h-1,k+\lfloor \Phi_p \rfloor-1}$. Note that because $g_{i,j}$ ends with node $v_{i+1,j-\lfloor \Phi_p \rfloor+1}$ and $k + \lfloor \Phi_p \rfloor - 1 \geq j - \lfloor \Phi_p \rfloor + 2$, then every path begin with $g_{i,j}$ and end with $g_{h-1,k+\lfloor \Phi_p \rfloor-1}$ must traverse an edge $e_{m,\ell}$ with $m \in [N]$ for each $\ell \in [j - \lfloor \Phi_p \rfloor + 2, k + \lfloor \Phi_p \rfloor - 1]$. We arbitrarily picked $m = 1$ for each ℓ, then we can conclude that

$$w(\bar{P}^1_{v_1,v_2}) \geq \sum_{\ell=j-\lfloor \Phi_p \rfloor+2}^{k+\lfloor \Phi_p \rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1). \tag{20}$$

Then according to the first inequality of Claim 1, we have

$$\sum_{\ell=i}^{h} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h} w(P^1_\ell). \tag{21}$$

Because $v_1 \in P^1_i$ and $v_2 \in P^1_h$, so $P^1_{v_1,v_2}$ contains subpath $P^1_{i+1}, \dots, P^1_{h-1}$ and part of subpaths P^1_i and P^1_h. So we can conclude that

$$\sum_{\ell=i}^{h} \text{cost}(P^1_\ell) \geq w(P^1_{v_1,v_2}) \tag{22}$$

By combining inequalities (20), (21) and (22) we have $w(\bar{P}^1_{v_1,v_2}) \geq w(P^1_{v_1,v_2})$. □

3.2 Upper bound on the PoA in series-parallel networks

In this section, we prove the upper bound on the PoA stated in Theorem 1.2. 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 [13]. 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 Section 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 3.3. *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. □

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 3.4. *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

$$\sum_{e \in p}(o_e+1)d_e(o_e+1) - \sum_{e \in p} o_e d_e(o_e) < \frac{\text{cost}(f)}{N}. \tag{23}$$

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

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

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

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

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

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

By Lemma 3.3, since C_i^+ is an (s,t)-path, we have $\mathrm{cost}_o(C_i^+) = \sum_{e \in C_i^+} d_e(o_e) \geq \frac{\mathrm{cost}(f)}{N}$. Thus, by (23) we obtain $\mathrm{cost}(o') - \mathrm{cost}(o) < 0$, which contradicts the fact that o is a social optimum. □

By using Lemma 3.4, we can derive a lower bound on $\mathrm{cost}_o(p)$ similar to the lower bound on $\mathrm{cost}_o(C_i^+)$ stated in Lemma 3.3, but with an extra factor of $y(\mathcal{D})$.

Lemma 3.5. *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 $p \in P$,*

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

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

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

where the first inequality follows the definition of $y(\mathcal{D})$ stated in Equation (1) and the second inequality follows from Lemma 3.4. □

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

Proof. By Lemma 3.5 we know that given an arbitrary decomposition P of the social optimal flow o, for all $p \in P$, we have $y(\mathcal{D}) \operatorname{cost}_o(p) \geq \frac{\operatorname{cost}(f)}{N}$. Then we can conclude that:

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

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

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 [13]. For each $i \in [h]$, we identify the pair of nodes u_i, v_i and we define

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

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

Lemma 3.7. *If G is series-parallel, then $\mathcal{L} = \{E_1, \ldots, E_h\}$ 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. □

By Proposition 1 in [13], 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 3.8. *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 $\{C_1, \ldots, C_h\}$ 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 3.7 $\mathcal{L}$ 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. □

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

Lemma 3.9. *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 3.8, 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 3.6 we conclude that the lemma holds. □

We are finally ready to prove Theorem 1.2, 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 1.2. 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 $E_{C_1}, \ldots, E_{C_l}$ 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, $E_{C_1}, \ldots, E_{C_l}$ 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

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

where the inequality follows from the fact that $y(\mathcal{D}) \geq 1$ and from Lemma 3.9. □

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 $j = 0, \ldots, p$.

Lemma 3.10. *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 $x \in \mathbb{N}^+$

$$\begin{aligned} y(\text{Poly-}p) &= \sup_{a_0,\ldots,a_p \in \mathbb{R}_{\geq 0},\ x \in \mathbb{N}^+} \frac{(x+1)\sum_{j=0}^{p} a_j (x+1)^j - x \sum_{j=0}^{p} a_j x^j}{\sum_{j=0}^{p} a_j x^j} \\ &= \sup_{a_0,\ldots,a_p \in \mathbb{R}_{\geq 0},\ x \in \mathbb{N}^+} \frac{\sum_{j=0}^{p} \left(a_j \left((x+1)^{j+1} - x^{j+1}\right)\right)}{\sum_{j=0}^{p} a_j x^j}. \end{aligned} \tag{24}$$

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

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

As a consequence, we can upper bound (24) by

$$\max_{j \in \{0,\ldots,p\}, x \in \mathbb{N}^+} \frac{(x+1)^{j+1} - x^{j+1}}{x^j}. \tag{25}$$

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 (25) we then obtain

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

□

By Theorem 1.2 and Lemma 3.10 we obtain that the PoA of series-parallel network congestion games with polynomial delay functions with highest degree is p is at most $2^{p+1} - 1$.

3.3 Lower bound on the PoA in series-parallel networks

In this section, we illustrate how to construct a family of instances that asymptotically achieve the lower bound on the PoA stated in Theorem 1.3. 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, $s_1 = \cdots = s_m = \frac{\sum_1^m q_i}{m}$ 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 1.3, 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 1.3 is given in the appendix.

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

$$1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r = 1 - \frac{1}{4} r^{2-\frac{1}{2^p}} + r - r^{\frac{1}{2^p}}. \tag{26}$$

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 $\lim_{p\to\infty} \frac{1-\left(1-\frac{1}{2^p}\right)\left(\frac{2}{2^{p+1}-1}\right)^{\frac{1}{2^p-1}}}{\frac{p}{2^p}} = 0$, proving that (26) is in $O\left(p/2^p\right)$, which implies that when N goes to infinity the PoA is at least in $\Omega\left(2^p/p\right)$.

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, $\mathrm{tot}(P)$ 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 1.4.

Proof of Theorem 1.4. Let P_o be the social optimum with respect to the total cost, and let $P_{\hat{o}}$ 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 $\max(P_f) \leq z(\mathcal{D})y(\mathcal{D})\max(P_{\hat{o}})$.

Because P_f is a PNE and $\max(P_f)$ is the cost of a player, we have $\max(P_f) \leq \mathrm{cost}^+_f(p^i_f)$ for any $i \in [N]$. Moreover, by (3), we have $\mathrm{cost}^+_f(p^i_f) \leq z(\mathcal{D})\,\mathrm{cost}_f(p^i_f)$. 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

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

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,

$$\begin{aligned} z(\mathcal{D})\,\mathrm{tot}(P_f) &\leq z(\mathcal{D})y(\mathcal{D})\,\mathrm{tot}(P_o) && (27)\\ &\leq z(\mathcal{D})y(\mathcal{D})\,\mathrm{tot}(P_{\hat{o}}) && (28)\\ &\leq z(\mathcal{D})y(\mathcal{D})(N \cdot \max(P_{\hat{o}})). && (29) \end{aligned}$$

Inequality (27) directly follows Theorem 1.2. Inequality (28) holds since P_o is the social optimum state with respect to the total cost, which implies that $\mathrm{tot}(P_o) \leq \mathrm{tot}(P_{\hat{o}})$. Inequality (29) holds because $\max(P_{\hat{o}})$ is the maximum player's cost in $P_{\hat{o}}$. ☐

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 4.1. *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 $b_0, \ldots, b_p$ and $c_0, \ldots, c_p$, we have

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

Thus,

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

☐

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 1.5. 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 $\mathrm{cost}(P_f)/\mathrm{cost}(P_o)$. 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

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

We construct G' as follows. First, let $G_1, \ldots, G_N$ be N duplicates of G and let G' be the (s,t)-network obtained by composing in series $G_1, \ldots, G_N$. We remark that any graph structure possessed by G is still valid for G', by our assumption. Let $P_f = \{p_f^1, \ldots, p_f^N\}$ and $P_o = \{p_o^1, \ldots, p_o^N\}$. For each $i \in [N]$ let $P_{f_i} = \{p_{f_i}^1, \ldots, p_{f_i}^N\}$ and $P_{o_i} = \{p_{o_i}^1, \ldots, p_{o_i}^N\}$ 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 $j(i) = (i + N - 1) \mod N$. For example, the strategy of player 2 in $P_{f'}$ is obtained by composing in series the paths $p_{f_1}^2, p_{f_2}^3, \ldots, p_{f_{N-1}}^N, p_{f_N}^1$. 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 $P_{f'} = \{p_{f'}^1, \ldots, p_{f'}^N\}$ 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 $P_{o'} = \{p_{o'}^1, \ldots, p_{o'}^N\}$ 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.

□

5 Conclusion

Our contributions fill a gap in the literature on the PoA of atomic, unweighted, symmetric network congestion games. We have investigated the impact of both symmetry and of network structure on the worst-case PoA in network congestion games. Previous works had either addressed asymmetric games over general networks, or symmetric games over very simple network structures, such as parallel-link networks and extension-parallel networks. First, we considered symmetric network congestion games over arbitrary networks and showed that when the delay functions are polynomial, the worst-case PoA is very close to that of asymmetric network congestion games. This implies that symmetry does not influence the worst-case PoA significantly.

Then, we considered the class of series-parallel networks, corresponding to graphs with treewidth 2. These networks arise in many applications and understanding how their structure impacts the PoA in network congestion games could be the first step towards relating the worst-case PoA to the treewidth parameter. Our results indicates that, when restricting from symmetric games over arbitrary networks to symmetric games over series-parallel networks, the worst-case PoA significantly drops. On the other hand, the worst-case PoA quickly degrades when going from extension-parallel to series-parallel networks.

In this paper we have focused on symmetric games, but it is not clear if network structure could affect the PoA in asymmetric games. An open question is: in asymmetric (unweighted) congestion games over other special network structures, for example series-parallel networks, is the PoA still significantly smaller than that over arbitrary networks?

References

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *STACS 2006*, pages 218–229, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[2] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '05, page 57–66, New York, NY, USA, 2005. Association for Computing Machinery.

[3] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. *SIAM Journal on Computing*, 42(1):160–177, 2013.

[4] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In Mark de Berg and Ulrich Meyer, editors, *Algorithms – ESA 2010*, pages 17–28, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[5] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Econ. Comput.*, 2(4), October 2014.

[6] V. Bilò and C. Vinci. On the Impact of Singleton Strategies in Congestion Games. In Kirk Pruhs and Christian Sohler, editors, *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:14, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[7] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, nov 2011.

[8] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '05, page 67–73, New York, NY, USA, 2005. Association for Computing Machinery.

[9] J. Correa, J. de Jong, B. de Keijzer, and M. Uetz. The inefficiency of Nash and subgame perfect equilibria for network routing. *Mathematics of Operations Research*, 44(4):1286–1303, November 2019.

[10] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Trans. Algorithms*, 3(1), feb 2007.

[11] A. Epstein, M. Feldman, and Y. Mansour. Efficient graph topologies in network routing games. *Games Econ. Behav.*, 66(1):115–125, 2009.

[12] D. Fotakis. Stackelberg strategies for atomic congestion games. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *Algorithms – ESA 2007*, pages 299–310, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[13] D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. In Burkhard Monien and Ulf-Peter Schroeder, editors, *Algorithmic Game Theory*, pages 33–45, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[14] D. Fotakis. Stackelberg strategies for atomic congestion games. *Theor. Comp. Sys.*, 47(1):218–249, jul 2010.

[15] D. Fotakis and G. P. Spirakis. Cost-balancing tolls for atomic network congestion games. *Internet Mathematics*, 5(4):343–363, 2008.

[16] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. In *Automata, Languages and Programming*, pages 645–657, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[17] M. Gairing and F. Schoppmann. Total latency in singleton congestion games. In *Internet and Network Economics. WINE 2007*, volume 4858 of *Lecture Notes in Computer Science*, pages 381–387, Berlin, Heidelberg, 2007. Springer.

[18] B. Hao and C. Michini. The price of anarchy in series-parallel network congestion games. *Mathematical Programming*, pages 1–31, 2022.

[19] P. T. Harker. Multiple equilibrium behaviors on networks. *Transportation Science*, 22(1):39–46, 1988.

[20] R. Holzman and D. Monderer. Strong equilibrium in network congestion games: Increasing versus decreasing costs. *Int. J. Game Theory*, 44(3):647–666, aug 2015.

[21] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.

[22] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *STACS 2004*, volume 2996 of *Lecture Notes in Computer Science*, pages 547–558, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[23] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406(3):187–206, 2008. Algorithmic Aspects of Global Computing.

[24] J. Nickerl. The minimum tollbooth problem in atomic network congestion games with unsplittable flows. *Theor. Comp. Sys.*, 65(7):1094–1109, oct 2021.

[25] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

[26] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[27] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341 – 364, 2003. Special Issue on STOC 2002.

[28] T. Roughgarden. The maximum latency of selfish routing. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '04, page 980–981, USA, 2004. Society for Industrial and Applied Mathematics.

[29] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 513–522, New York, NY, USA, 2009. Association for Computing Machinery.

[30] T. Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5), November 2015.

[31] T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, March 2002.

[32] S. Suri, C. D. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. In *Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '04, page 188–195, New York, NY, USA, 2004. Association for Computing Machinery.

[33] J. Valdes, R. Tarjan, and E. Lawler. The recognition of series parallel digraphs. *SIAM Journal on Computing*, 11(2):298–16, 05 1982.

Appendix

Proof of Theorem 1.3

We here provide the complete proof of Theorem 1.3. Consider the construction described in Section 3.3, which is represented in Fig.2. We will first show that this construction satisfies the properties stated in the next two lemmas.

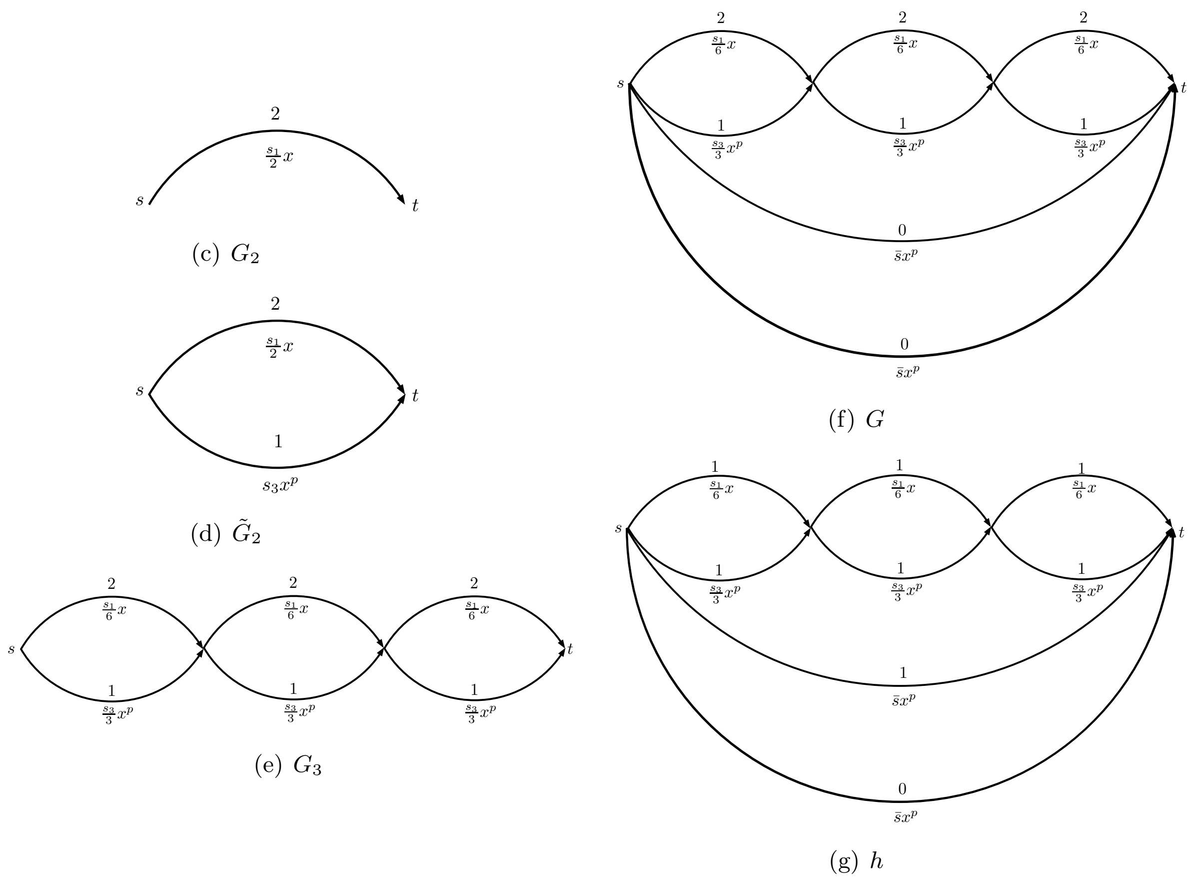

Figure 2: Given an input sequence $\{q_1, q_2, q_3\}$, $p \in \mathbb{N}^+$ and $m = 2$, we first average the first m numbers and get $\{s_1, s_2, s_3\}$, where $s_1 = s_2 = \frac{q_1+q_2}{2}$, $s_3 = q_3$ and $\bar{s} = \frac{s_1+s_2+s_3}{3}$. (2(f)) is the output network G and its corresponding PNE flow f. (2(c)), (2(d)), (2(e)) are the intermediate networks and flows according to our construction. (2(g)) is the flow h defined in the proof of Theorem 1.3 where $k = 1$.

Lemma .1. *The (s,t)-flow f has an (s,t)-path $\bar{p}$ with flow value m and* $\text{cost}_f(\bar{p}) = s_1$.

Proof. We prove the lemma by induction on $i \in \{m, \dots, N\}$. The base case is $i = m$. In this case $f_i = f_m$ is a flow of value m on a single (s,t)-edge with delay function $\frac{s_1 x}{m}$. The path $\bar{p}^m$ defined by this edge has cost $\text{cost}_{f_m}(\bar{p}^m) = s_1$.

Suppose that for each $m \le i < N$ it holds that f_i has an (s,t)-path $\bar{p}^i$ with flow value m and $\text{cost}_{f_i}(\bar{p}^i) = s_1$. We first construct $\tilde{f}_i$ by composing in parallel f_i and a new (s,t)-edge. Clearly, $\bar{p}^i$ has still flow value m and $\text{cost}_{\tilde{f}_i}(\bar{p}^i) = s_1$. Then we compose in series $i+1$ copies of flow $\tilde{f}_i$ to get f_{i+1} and we divide the delay functions by $i+1$. The new (s,t)-path $\bar{p}^{i+1}$ is obtained by composing in series $i+1$ copies of $\bar{p}^i$. By construction, this path has flow value m and $\text{cost}_{f_{i+1}}(\bar{p}^{i+1}) = s_1$. $\square$

Lemma .2. *The (s,t)-flow f has cost 1, and it can be decomposed into N (s,t)-paths $\{p^1, \ldots, p^N\}$ that define a PNE in G. Moreover $\text{cost}_f(p^i) = 1/N$ for all $i \in [N]$, i.e., each player incurs the same cost.*

Proof. First, we show that f_N has cost $\sum_{i=1}^N s_i = 1$ and it can be decomposed into a PNE in G_N where each player incurs the same cost. We show this by induction on i. When $i = m$, G_m is a single (s,t)-edge, and f_m is an (s,t)-flow of value m routed through this edge. Moreover, $\text{cost}(f_m) = \frac{s_1 m}{m} m = \sum_{i=1}^m s_i$. Note that we cannot define any alternative flow in G_m. Moreover, f_m admits a unique decomposition into N (s,t)-paths, thus f_m is a PNE flow where each player uses the same edge and incurs the same cost.

Now we assume that when $i = k$, f_k has cost $\sum_{i=1}^k s_i$, and it can be decomposed into a PNE in G_k where each player incurs the same cost. Our goal is to prove that the same holds for $i = k+1$. Note that in our construction first we define $\tilde{G}_k$ and $\tilde{f}_k$ by composing in parallel f_k and a new (s,t)-edge with delay $s_{k+1}x$ and flow value 1. Thus, we first show that $\tilde{f}_k$ is a PNE flow in $\tilde{G}_k$. By the inductive hypothesis, flow f_k can be decomposed into a PNE in G_k where each player's cost is $\frac{1}{k}\sum_{i=1}^k s_i$. To define a decomposition of $\tilde{f}_k$, we augment the decomposition of f_k by appending the extra (s,t)-edge used to construct $\tilde{G}_k$. Clearly, $\text{cost}(\tilde{f}_k) = \text{cost}(f_k) + s_{k+1} = \sum_{i=1}^{k+1} s_i$. Moreover, (i) no player paying $\frac{1}{k}\sum_{i=1}^k s_i$ has an incentive to deviate, since $2^p s_{k+1} \geq \frac{1}{k}\sum_{i=1}^k s_i$, and (ii) the player paying s_{k+1} does not deviate since s_{k+1} is the minimum cost (s,t)-path in $\tilde{f}_k$. This shows that $\tilde{f}_k$ is a PNE flow in $\tilde{G}_k$. Recall that in our construction we define G_{k+1} and f_{k+1} by composing in series $k+1$ copies of $\tilde{G}_k$ with flow $\tilde{f}_k$, and we divide all the delay functions by $k+1$. Clearly, $\text{cost}(f_{k+1}) = \text{cost}(\tilde{f}_k) = \sum_{i=1}^{k+1} s_i$. We define a decomposition of f_{k+1} into $k+1$ (s,t)-paths as follows. Since there are $k+1$ players and $k+1$ identical copies of $\tilde{G}_k$ composed in series, we let each player choose their original strategy in f_k in k components, and choose the extra edge used to define $\tilde{G}_k$ in one component. Thus, in this decomposition of f_{k+1} each player incurs the same cost and no player has an incentive to deviate from their strategy.

Finally, we show that $f = f_N$ is a PNE flow on G and $\text{cost}(f) = \text{cost}(f_N) = \sum_1^N s_i = 1$. Recall that we construct G by composing in parallel G_N and m new (s,t)-edges $e_1, \ldots, e_m$ with delay function $\frac{1}{N}x^p$. Since in f every player incurs a cost equal to $\frac{1}{N}$, no player has an incentive to deviate to an edge e_i, $i \in [m]$. Thus, f is a PNE flow on G.

□

Define $\mu(m,N) = \prod_{j=m}^{N-1} \frac{2^p j}{2^p j+1}$. We will need the results stated in the next two lemmas.

Lemma .3. *Let $\{q_1, \ldots, q_N\}$ 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 $i \in [N-1]$. Then for every $m \in [N]$ we have $\sum_{i=1}^m q_i = \mu(m,N)$.*

Proof. We proceed by induction on m. The base case is $m = N-1$. Since $q_N = \frac{1}{2^p(N-1)}\sum_{j=1}^{N-1} q_j$, we have:

$$\sum_{j=1}^{N-1} q_j = 1 - q_N = 1 - \frac{1}{2^p(N-1)}\sum_{j=1}^{N-1} q_j. \tag{30}$$

By equation (30), we have $\frac{2^p(N-1)+1}{2^p(N-1)}\sum_{j=1}^{N-1} q_j \leq 1$. This implies that $\sum_{j=1}^{N-1} q_j \leq \frac{2^p(N-1)}{2^p(N-1)+1} = \mu(N-1,N)$. Thus the statement holds for the base case.

Next we assume that the statement holds for $m \in \{k, \ldots, N-1\}$, and we prove that it also holds for $m = k-1$. Based on our inductive hypothesis, $\sum_{j=1}^k q_j \leq \mu(k,N)$. Moreover, since

$q_k = \frac{1}{2^p(k-1)} \sum_{j=1}^{k-1} q_j$, we have:

$$\sum_{j=1}^{k-1} q_j = \sum_{j=1}^{k} q_j - q_k = \mu(k,N) - \frac{1}{2^p(k-1)} \sum_{j=1}^{k-1} q_j. \tag{31}$$

According to (31), we have $\frac{2^p(k-1)+1}{2^p(k-1)} \sum_{j=1}^{k-1} q_j = \mu(k,N)$. This implies that $\sum_{j=1}^{k-1} q_j = \frac{2^p(k-1)}{2^p(k-1)+1}\mu(k,N) = \mu(k-1,N)$. Thus, the statement holds. □

Lemma .4. *For $m \in [N-1]$ we have $\sqrt[2^p]{\frac{2^p m-(2^p-1)}{2^p N-(2^p-1)}} \leq \mu(m,N)$.*

Proof. First we can equivalently write:

$$\mu(m,N) = \prod_{j=m}^{N-1} \frac{2^p j}{2^p+1} = \sqrt[2^p]{\left(\prod_{j=m}^{N-1} \frac{2^p j}{2^p j+1}\right)^{2^p}}.$$

We lower bound the argument of the square root as follows.

$$\begin{aligned}\prod_{j=m}^{N-1} \left(\frac{2^p j}{2^p j+1}\right)^{2^p} &\geq \prod_{j=m}^{N-1} \left(\prod_{k=0}^{2^p-1} \frac{2^p j-k}{2^p j+1-k}\right) \\ &= \frac{2^p m-(2^p-1)}{2^p N-(2^p-1)}.\end{aligned}$$

□

We will now use the results stated in the above lemmas to prove Theorem 1.3.

Proof of Theorem 1.3. Consider the network congestion game on the network G defined above. By Lemma .1, f has an (s,t)-path $\bar{p}$ with flow value m and $\text{cost}_f(\bar{p}) = s_1$. For each edge e in $\bar{p}$, let $a_e x^p$ be the delay function of e. Note that $\text{cost}_f(\bar{p}) = \sum_{e\in\bar{p}} a_e m = s_1$ implies that $\sum_{e\in\bar{p}} a_e = \frac{s_1}{m}$. Recall that $r = \frac{m}{n}$ and $l = \frac{k}{m}$. Define h as the flow obtained from f by moving a subflow of value $(m-k)$ from p to the (s,t)-edges $e_1, \ldots, e_{m-k}$, which have all delay function $\frac{1}{N}x^p$. Then by construction we have:

$$\begin{aligned}\text{cost}(f) - \text{cost}(h) &= m\,\text{cost}_f(\bar{p}) - \left(k\,\text{cost}_h(\bar{p}) + (m-k)\frac{1}{N}\right) \\ &= s_1 m - \left(\frac{s_1}{m}k^2 + (m-k)\frac{1}{N}\right) \qquad (32) \\ &= \left(\frac{s_1}{m}m^2 - \frac{s_1}{m}k^2 - \frac{m-k}{m}ms_1\right) + \frac{m-k}{m}\left(ms_1 - \frac{m}{N}\right) \\ &= \left(\frac{s_1}{m}mk - \frac{s_1}{m}k^2\right) + \frac{m-k}{m}\left(\sum_{i=1}^{m} s_i - \frac{m}{N}\right), \qquad (33)\end{aligned}$$

where equality (32) holds since $\sum_{e\in\bar{p}} a_e = \frac{s_1}{m}$. Equality (33) holds since the first m of s_i are equal.

By Lemma .3 and Lemma .4 we have

$$\sum_{i=1}^{m} q_i - \frac{m}{N} = \mu(m,N) - \frac{m}{N} \geq \left[(\sqrt[2^p]{r} - r) - \epsilon\right]. \tag{34}$$

Now observe that

$$\frac{s_1}{m}m^2 = ms_1 = \frac{m}{N} + (\sum_1^m s_i - \frac{m}{N}) \geq r + \left[(\sqrt[2^p]{r} - r) - \epsilon\right] = (\sqrt[2^p]{r} - \epsilon), \tag{35}$$

where the inequality follows from (34) and the fact that $\sum_1^m s_i = \sum_1^m q_i$.

This implies

$$\frac{s_1}{m}mk - \frac{s_1}{m}k^2 = (l - l^2)\frac{s_1}{m}m^2 \geq (l - l^2)(\sqrt[2^p]{r} - \epsilon), \tag{36}$$

where the inequality follows from (35).

From (33) and (36)we obtain

$$\begin{aligned} \mathrm{cost}(f) - \mathrm{cost}(h) &\geq (l - l^2)(\sqrt[2^p]{r} - \epsilon) + (1 - l)(\sum_1^m s_i - \frac{m}{N}) \\ &\geq (l - l^2)(\sqrt[2^p]{r} - \epsilon) + (1 - l)\left[(\sqrt[2^p]{r} - r) - \epsilon\right], \end{aligned} \tag{37}$$

where inequality (37) follows from (34). By Lemma .2 we know that $\mathrm{cost}(f) = \sum_1^N s_i = 1$, thus we obtain:

$$\begin{aligned} \mathrm{cost}(h) &\leq 1 - (l - l^2)(\sqrt[2^p]{r} - \epsilon) - (1 - l)\left[(\sqrt[2^p]{r} - r) - \epsilon\right] \\ &= 1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r + (1 - l^2)\epsilon \end{aligned} \tag{38}$$

To obtain an upper bound on $\mathrm{cost}(h)$ we minimize the right-hand-side of (38) with respect to r and l. Observe that $\epsilon \to 0$ when $N \to \infty$, thus we solve

$$\begin{aligned} \min\ & l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r \\ \text{s.t.}\ & r \in [0, 1), l \in [0, 1], \end{aligned}$$

which is achieved at $r = \left(\frac{2}{2^{p+1}-1}\right)^{\frac{2^p}{2^p-1}}$, $l = \frac{1}{2}r^{1-\frac{1}{2^p}}$. Since $\frac{\mathrm{cost}(f)}{\mathrm{cost}(o)} \geq \frac{\mathrm{cost}(f)}{\mathrm{cost}(h)}$, we obtain a lower bound for the PoA. □