The Price of Anarchy in Series-Parallel Network Congestion Games

Bainian Hao [*]   Carla Michini [†]

September 30, 2021

Abstract

We study the inefficiency of pure Nash equilibria in symmetric network congestion games defined over series-parallel networks with affine edge delays. For arbitrary networks, Correa (2019) proved a tight upper bound of 5/2 on the PoA. On the other hand, for extension-parallel networks, a subclass of series-parallel networks, Fotakis (2010) proved that the PoA is 4/3. He also showed that this bound is not valid for series-parallel networks by providing a simple construction with PoA 15/11. Our main result is that for series-parallel networks the PoA cannot be larger than 2, which improves on the bound of 5/2 valid for arbitrary networks. We also construct a class of instances with a lower bound on the PoA that asymptotically approaches 27/19, which improves on the lower bound of 15/11.

1 Introduction

Network congestion games are commonly used to model problems in large-scale networks such as routing in communication networks and traffic planning in road networks [16, 24]. In a network congestion game there is a finite number of selfish players, and each of them has to select a path from an origin to a destination. The edges of the network are regarded as resources that can get congested, because each player using an edge experiences a delay that is non-decreasing with respect to the total number of players using it. Each player aims at minimizing the cost of the path she selects, which is the sum of the delays of all the edges in the path.

A *pure Nash equilibrium* (PNE) is a configuration where no player can decrease her cost by unilaterally deviating to another path, and it represents a stable outcome of the game. However, since the players act selfishly and independently in a non-cooperative fashion, a PNE might be far from minimizing the social cost, which is the sum of all players' costs.

The inefficiency arising from this lack of coordination is quantified by two measures: the *Price of Anarchy* (PoA) and the *Price of Stability* (PoS). The PoA, introduced by

---

[*]Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI, USA. E-mail: `bhao8@wisc.edu`.

[†]Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI, USA. E-mail: `michini@wisc.edu`.

Koutsoupias and Papadimitriou [20], is the largest ratio between the cost of a PNE and the minimum social cost. The PoS, introduced by Anshelevich et al. [2, 3], is the smallest ratio between the cost of a PNE and the minimum social cost. Thus, the PoA and the PoS measure the inefficiency of a PNE in the worst-case and best-case scenarios, respectively.

The PoA and PoS of network congestion games have been studied in a number of settings [28, 12, 29]. In non-atomic games, players are infinitesimal agents that control only a negligible amount of flow and cannot affect each other, while in atomic games players control a larger, non-negligible amount of flow. Moreover, in the atomic case players may or may not be allowed to split their flow along different paths. Finally, special cases have been investigated, which arise from assuming some structural properties on the graph topology and/or on the delay functions [15, 17, 10, 7, 8]. In this paper, we consider the atomic setting, and we assume that each player controls one unit of flow that has to be routed on a single path. Moreover, we consider symmetric games, where all the players have the same origin-destination pair. We focus on the special case where the network is a two-terminal series-parallel graph and the edge delays are affine functions, see Fig. 1 for an example. First, two-terminal series-parallel networks can be recognized in linear-time [32] and are relevant in many applications, such as for problems on electric networks, scheduling and compiler optimization. Moreover, the special structure of these graphs and their decomposition properties can be exploited to define efficient algorithms for combinatorial problems that are NP-hard in general [6, 19, 31]. Finally, series-parallel graphs are graphs with treewidth 2, thus understanding how their structure impacts the PoA in network congestion games could be the first step towards relating the PoA to the treewidth parameter. Indeed, exploiting the structure of series-parallel networks is crucial to prove our main result.

Theorem 1 *Suppose that G is a series-parallel (s, t)-graph and that the delay functions are affine. Then the PoA is at most 2.*

The best upper bound on the PoA in series-parallel network congestion games with affine delays that was previously known was equal to $5/2$, however this bound actually holds for network congestion games on arbitrary graphs [10]. In contrast, for network congestion games with affine delays in extension-parallel graphs, Fotakis [15] proved a bound of $4/3$ on the PoA. We recall that extension-parallel networks, similarly to series-parallel networks, can also be obtained by parallel and series compositions of extension-parallel components, but in every series composition at least one component must be a single edge. Thus extension-parallel networks are a subclass of series-parallel networks, and indeed they display much stronger properties than series-parallel networks. Notably, paths in extension-parallel networks are *linearly independent* [22], in the sense that every path contains an edge not included in any other path. This property is crucially exploited by Fotakis to prove the bound of $4/3$ on the PoA. However, neither this property, nor the bound of $4/3$ on the PoA are valid for the larger class of series-parallel networks.

In fact, Fotakis provided a counterexample of a series-parallel network where the PoA is equal to $15/11 > 4/3$ [15]. This was the best lower bound on the PoA known so far for symmetric network congestion games on series-parallel networks and affine delays. We improve such lower bound by constructing a class of instances with a lower bound on the PoA that asymptotically approaches $27/19$ as the number of players goes to infinity.

Theorem 2 *The PoA of series-parallel congestion games with affine delays is at least $27/19 - \epsilon$, where $\epsilon \to 0$ as $N \to \infty$.*

Related Work. The complexity of finding a PNE in network congestion games and more general congestion games has been widely investigated in the literature. These games belong to the class of potential games, for which a PNE is guaranteed to exists. Potential games are characterized by the existence of a *potential function*, and each local optimum of such function corresponds to a PNE [26, 27, 23]. Fabrikant et al. [14] gave a strongly polynomial algorithm to find a PNE in symmetric network congestion games, and proved that in the asymmetric case network congestion games are PLS-complete, even in the case of linear delays [18, 14, 1].

The inefficiency of pure Nash equilibria in network congestion games and congestion games has also been widely studied. One way to measure such inefficiency is by comparing the social cost of a PNE and the minimum social cost. The Price of Anarchy (PoA) is the largest ratio between the cost of a PNE and the minimum social cost [20].

For non-atomic network congestion games with delay functions in class $\mathcal{D}$, Roughgarden proved a tight bound on the PoA that is independent from the network structure [30, 28]. The bound is a function $\rho(\mathcal{D})$ that only depends on $\mathcal{D}$ and is equal to $4/3$ for affine delays. Later, Correa et al. [11, 12] provided a unifying framework to study the PoA by introducing a function $\beta(\mathcal{D}) = 1 - 1/\rho(\mathcal{D})$.

For general atomic congestion games with affine delays and $N \geq 3$ players, Awerbuch et al. [4, 5] and Christodoulou and Koutsoupias [9] independently provided an upper bound of $5/2$ on the PoA. If the game is symmetric the bound can be improved to $(5N-2)/(2N+1)$ [9]. Correa et. al. [10] later proved this bound is *tight* for symmetric network congestion games with linear delays, by exhibiting a family of instances (parametrized by N) that achieves this bound. Each instance is composed by N disjoint (s,t)-paths, plus some *connecting* edges that link these paths. We remark that this construction inherently violates the structure of series-parallel networks, because of the presence of the connecting edges.

When assuming special properties on the network topology and/or on the delay functions the above bounds can be improved. A notable example are network congestion games on extension-parallel networks, a subclass of series-parallel networks The structure of these networks can be exploited to prove several nice properties of extension-parallel network congestion games [22, 13, 17, 15]. In particular, Fotakis studied the PoA in extension-parallel network congestion games and established that for delay functions in class $\mathcal{D}$ the PoA is at most $\rho(\mathcal{D})$ in the symmetric case [15]. For affine delay functions this bound is equal to $4/3$. However the bound is not valid for the larger class of series-parallel networks. In fact, Fotakis provided a counterexample of a series-parallel network where the PoA is equal to $15/11 > 4/3$ [15], see Fig. 1. Further structural properties have been investigated in the literature. For example, Bilò and Vinci [8] improved the bounds of [4, 5, 9, 15] by making the the additional assumption that the costs of any two strategies available to a same player, when evaluated in absence of congestion, are within a factor θ one from the other. Bhaskar et al. [7] instead focused on the price of collusion in atomic splittable congestion games on series-parallel networks and proved that the PoA cannot exceed the number of players.

Our approach. To prove that in series-parallel network congestion games with affine delays the PoA is at most 2, we need to overcome some of the limitations in Fotakis' approach [15], which is tailored to extension-parallel networks.

The first crucial property exploited in [15] is that there is a one-to-one correspondence between strategy profiles and network flows. Specifically, for a game with N players having origin s and destination t, each (s, t)-flow of value N corresponds to a unique strategy profile (up to players' permutation), because there is a unique decomposition of the (s, t)-flow into N (s, t)-paths. The second crucial property exploited in [15] is that all pure Nash equilibria are *global* optima of the potential function. This can be used to show that the PoS and the PoA coincide.

Both properties do *not* extend to series-parallel networks. In particular, each (s, t)-flow f of value N can be decomposed into different strategy profiles, and while some of them might be a PNE, some of them might not. We define a *greedy decomposition* of f into the single players' strategies. A similar definition was introduced to compute to generalized maximum flows in series–parallel graphs [21, 25]. We prove several properties of greedy decompositions and we crucially exploit these properties to derive to derive a bound on the PoA. .

For a PNE (s, t)-flow f and a (s, t)-flow o minimizing the total cost, we use the function $\Delta(f, o)$ defined in [15] to measure "how different" these two (s, t)-flows are. Fotakis proved that $\Delta(f, o) \leq 0$ if and only if f minimizes the potential function, and showed how this implies that the ratio between the cost of f and the cost of o is at most $4/3$. However, series-parallel networks might admit a PNE that is only a *local* optimum of the potential function. The crucial block of our proof consists in establishing that for any PNE f we have $\Delta(f, o)$ is at most $\frac{1}{4}$ the cost of f, which will imply that the PoA is at most 2.

2 Preliminaries

Notation. For a network G we denote by $V(G)$ and $E(G)$ the node set and the edge set of G, respectively. An edge $e \in E(G)$ can be explicitly written as the ordered pair (u, v), where u is the *tail* of e and v is the *head* of e. Directed paths will be simply referred to as paths. Unless otherwise specified, we will only consider *simple* paths, i.e., paths that do not traverse any node multiple times. A path from node u to node v is called a (u, v)-path. We say that two (u, v)-paths in G are *internally disjoint* if they only intersect at u and v. Paths and cycles of G are regarded as sequences of edges, thus we may for example write $e \in p$ rather than $e \in E(p)$ for a path p.

Let $G = (V, E)$ be an (s, t)-network, i.e., a network with source s and sink t, and let $c \in \mathbb{R}^{E(G)}$. 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 might simply use the term *flow*, if the source and sink of the flow are clear from the context or not relevant for the discussion. For a path p in G we define $c(p) = \sum_{e \in p} c_e$, and for a flow f in G we define $c(f) = \sum_{e \in E(G)} c_e f_e$. Finally, for a vector $f \in \mathbb{R}^{E(G)}$ we define $E(f) = \{(u, v) : (e = (u, v) \in E(G)$ and $f_e > 0)$ or $(e = (v, u) \in E(G)$ and $f_e < 0)\}$. Correspondingly, we denote by $G(f)$ the network $(V(G), E(f))$. Note

4

that if $f \geq 0$, then $G(f)$ is a subgraph of G. Given two subsets A and B of $E(G)$, we denote by $A \Delta B = (A \setminus B) \cup (B \setminus A)$ the *symmetric difference* of A and B. 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 $i \in [N]$. On the other hand, each (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. The delay functions are called *affine* if $d_e(x) = a_e x + b_e$ with $a_e \geq 0, b_e \geq 0$ for every $e \in E$. We denote the cost of a path p in G with respect to a flow f by $\mathrm{cost}_f(p) = \sum_{e \in p} d_e(f_e)$. Thus, the cost incurred by player i in state P is $\mathrm{cost}_f(p^i)$. We also define $\mathrm{cost}_f^+(p) = \sum_{e \in p} d_e(f_e + 1)$. Finally, the cost of flow f in G is denoted by $\mathrm{cost}(f) = \sum_{e \in E} f_e d_e(f_e)$, and it corresponds to the sum of all players' costs. Note that $\mathrm{cost}_f(p)$ and $\mathrm{cost}(f)$ coincide with $c(p)$ and $c(f)$ if $c \in \mathbb{R}^E$ is defined as $c_e = d_e(f_e)$ for all $e \in E$.

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

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

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*, defined as a state that minimizes the *total cost* $\mathrm{cost}(f(P)) = \sum_{i \in [N]} \mathrm{cost}_{f(P)}(p^i)$ over all the states $P = (p^1, \ldots, p^N) \in X$. With a slight abuse of terminology, we say that an (s, t)-flow o is a social optimum if o minimizes $\mathrm{cost}(g)$ over all integral (s, t)-flows g of value N.

Price of Anarchy and Price of Stability. The *Price of Anarchy* (PoA) is the maximum ratio $\frac{\mathrm{cost}(f)}{\mathrm{cost}(o)}$ such that o is a social optimum and f is a PNE. In other words, to compute the PoA we consider the "worst" PNE, i.e., a PNE whose total cost is as large as possible. The *Price of Stability* (PoS) is the minimum ratio $\frac{\mathrm{cost}(f)}{\mathrm{cost}(o)}$ such that o is a social optimum and f is a PNE. In other words, to compute the PoS we consider the "best" PNE, i.e., a PNE whose total cost is as small as possible.

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.
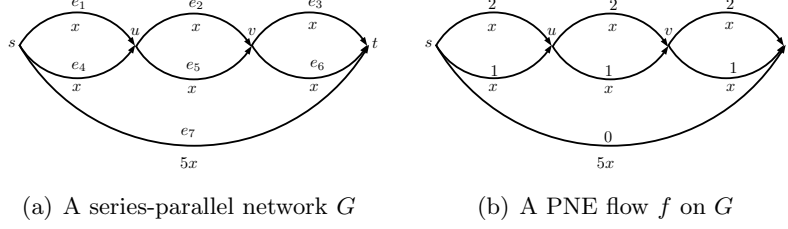
(a) A series-parallel network G (b) A PNE flow f on G

Figure 1: The series-parallel network congestion game of Example 1. The PoA is $15/11$.

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. Note that, if G is a series-parallel (s,t)-network, then every flow f of G is acyclic, i.e., there is no directed cycle in $E(f)$.

Example 1 *Consider the 3-player series-parallel congestion game with affine delays depicted in Fig. 1. The underlying network G and affine delay functions are showed in Fig. 1(a). A PNE flow f is represented in Fig. 1(b). The players' strategies are the (s,t)-paths $p^1 = (s,e_1,u,e_2,v,e_6,t)$, $p^2 = (s,e_1,u,e_5,v,e_3,t)$, and $p^3 = (s,e_4,u,e_2,v,e_3,t)$. Moreover, the flow o that minimizes the social cost can be reached from f by deviating one unit of flow from the path (s,e_1,u,e_2,v,e_3,t) to the path (s,e_7,t). As a result, we have $\mathrm{cost}(f) = 15$ and $\mathrm{cost}(o) = 11$, and this implies that the PoA is $15/11$. This example was originally introduced by Fotakis to show that the PoA of series-parallel congestion games with affine delays can be greater that $4/3$ [15].*

To prove Theorem 1, we will use the function $\beta(\mathcal{D}) := \sup_{d \in \mathcal{D}} \beta(d)$ introduced in [11], where $\mathcal{D}$ is a non-empty class of non-negative and non-decreasing functions and, for a non-negative and non-decreasing function $d(x)$, $\beta(d) := \sup_{x \geq y \geq 0} \frac{y(d(x)-d(y))}{xd(x)}$. We remark that when $\mathcal{D}$ is the class of affine functions, we have $\beta(\mathcal{D}) = 1/4$. Given an arbitrary PNE flow f of G and a social optimal flow o define

$$\Delta(f,o) := \sum_{e:f_e > o_e} (f_e - o_e)d_e(f_e) - \sum_{e:f_e < o_e} (o_e - f_e)d_e(f_e + 1).$$

By exploiting the definition of $\beta(\mathcal{D})$ the following inequality can be easily derived (see proof of Lemma 3 in [15]):

$$\mathrm{cost}(f) \leq \mathrm{cost}(o) + \beta(\mathcal{D})\mathrm{cost}(f) + \Delta(f,o). \tag{1}$$

If f is a global minimum of the potential function, then $\Delta(f,o) \leq 0$ [15]. However, series-parallel networks might admit PNE that do *not* minimize the potential function. To prove Theorem 1, we will exploit the special structure of series-parallel networks and affine delays, in order to show that $\Delta(f,o) \leq \frac{1}{4}\mathrm{cost}(f)$, see Theorem 3 and Corollary 4.

This immediately implies that $\text{cost}(f) \leq 2\text{cost}(o)$, establishing that the PoA is at most 2 for the case under consideration. The main ideas of the proof are described in the next Section.

3 Proof of Theorem 1

To prove Theorem 1 we use the following key result.

Theorem 3 *Suppose that G is a series-parallel (s,t)-network and that the delay functions are affine. Let f be a PNE and let $\mathcal{C} = \{(C_i^-, C_i^+) : i \in [k]\}$ be a collection of k (not necessarily distinct) pairs of internally disjoint (u_i, v_i)-paths in G, such that $|\{C_i^- : e \in C_i^-\}| \leq f_e$ for all $e \in E(G)$. Then*

$$\Delta(\mathcal{C}, f) := \sum_{i=1}^{k}(\text{cost}_f(C_i^-) - \text{cost}_f^+(C_i^+)) \leq \frac{1}{4}\text{cost}(f).$$

We will consider the network $G(o-f)$, which is a collection of simple cycles $\{C_1, \ldots, C_h\}$ 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 $C_i^- = \{e = (u,v) \in E : (v,u) \in C_i, o_e < f_e\}$. Since G is series-parallel, it is known that C_i^+ and C_i^- are two internally disjoint (u_i, v_i)-paths in G [15]. By defining $\mathcal{C}$ as the set containing s_i copies of (C_i^+, C_i^-) for each $i \in [h]$, we can apply Theorem 3. Since $\Delta(f, o)$ can be rewritten as

$$\Delta(f, o) = \sum_{i=1}^{h} s_i(\text{cost}_f(C_i^-) - \text{cost}_f^+(C_i^+)),$$

we have $\Delta(f, o) = \Delta(\mathcal{C}, f)$, and we obtain the following result.

Corollary 4 *Suppose that G is a series-parallel (s,t)-network and that the delay functions are affine. Let f be an arbitrary PNE and let o be a social optimum. Then $\Delta(f, o) \leq \frac{1}{4}\text{cost}(f)$.*

We remark that, to prove Theorem 1, it is sufficient to use inequality (1) in conjunction with Corollary 4.

4 Proof of Theorem 3

In this section, we formally prove Theorem 3. To this purpose, we need to first prove a number of intermediate results.

An acyclic (s,t)-flow f of value N can be decomposed into N simple (s,t)-paths in multiple ways. Given edge costs c_e, $e \in E$, we compute a c-greedy decomposition $\bar{P} = \{\bar{p}^1, \ldots, \bar{p}^N\}$ of f as follows. Set $f_1 = f$, $E_1 = E(f_1)$. At each step, compute the (s,t)-path $\bar{p}^i$ in (V, E_i) with highest cost with respect to c, and decrease the (s,t)-flow f_i by 1 on all the edges that belong to $\bar{p}^i$ to define f_{i+1} and E_{i+1}.

7

Example 1, continued *Consider again the network congestion game in Fig. 1 and its PNE flow f. We define the edge costs as $c_e = d_e(f_e)$, $e \in E(G)$. The c-greedy decomposition $\bar{P}$ of f consists of the (s,t)-paths $\bar{p}^1 = \bar{p}^2 = (s, e_1, u, e_2, v, e_3)$ and $\bar{p}^3 = (s, e_4, u, e_5, v, e_6)$. The costs of the paths $\bar{p}^1, \bar{p}^2, \bar{p}^3$ are $6, 6, 3$ respectively.*

In the next lemma, we prove a first basic property of greedy decompositions.

Lemma 5 *Let $c \in \mathbb{R}^E$ and suppose that G is a series-parallel (s,t)-network. For an (s,t)-flow f of G of value N, let P be an arbitrary decomposition of f into N (s,t)-paths $\{p^1, \ldots, p^N\}$ with $c(p^i) \geq c(p^{i+1})$, $i \in [N-1]$, and let $\bar{P} = \{\bar{p}^1, \ldots, \bar{p}^N\}$ be a c-greedy decomposition of f. Then $c(p^1) \leq c(\bar{p}^1)$ and $c(p^N) \geq c(\bar{p}^N)$.*

Proof. By construction, we have that $c(p^1) \leq c(\bar{p}^1)$.

We now prove $c(p^N) \geq c(\bar{p}^N)$ by proving that $\bar{p}^N$ is the cheapest path in $G(f)$ with respect to c. We proceed by induction on the number of edges $|F|$ in the network $G(f)$. If $|F| = 1$, then $G(f)$ is an edge. Thus we have $\bar{p}^1 = \cdots = \bar{p}^N$. This implies that $\bar{p}^N$ is the cheapest path.

Now we assume that when $|F| \leq k$, $\bar{p}^N$ is the cheapest path in f. When $|F| = k+1$, we have that f is composed by two flows f_1 and f_2 either in series or in parallel. Note that the number of edges $|F_1|$ in $G(f_1)$ and the number of edges $|F_2|$ in $G(f_2)$ are both at most k.

If f is composed in series by f_1, f_2, we can define from $\bar{P}$ two c-greedy decompositions $\bar{P}_1 = \{\bar{p}_1^1, \ldots, \bar{p}_1^N\}$, $\bar{P}_2 = \{\bar{p}_2^1, \ldots, \bar{p}_2^N\}$ of f_1 and f_2, respectively, such that $\bar{p}^i = \bar{p}_1^i \circ \bar{p}_2^i$ for all $i \in [N]$. By our inductive hypothesis, we have that $\bar{p}_1^N$ and $\bar{p}_2^N$ are the cheapest paths with respect to c in $G(f_1)$ and $G(f_2)$ respectively, thus $\bar{p}^N = \bar{p}_1^N \circ \bar{p}_2^N$ is the cheapest path with respect to c in $G(f)$.

If f is composed in parallel by f_1, f_2, we define $\bar{P}_1$ and $\bar{P}_2$ as the paths of $\bar{P}$ that belong to $G(f_1)$ and $G(f_2)$, respectively. Then $\bar{P}_1$ and $\bar{P}_2$ are c-greedy decompositions of f_1 and f_2, respectively. By our inductive hypothesis, the last path in $\bar{P}_1$ is the cheapest path with respect to c in $G(f_1)$, and the last path in $\bar{P}_2$ is the cheapest path with respect to c in $G(f_2)$. The cheapest among these two paths is the last path in $\bar{P}$, and it must be the cheapest path with respect to c in $G(f)$. $\square$

For a collection of N paths $P = \{p^1, \ldots, p^N\}$, $c \in \mathbb{R}^E$ and $x \geq 0$ we define $R(P, c, x) = \sum_{i=1}^{N} \max\{0, c(p^i) - x\}$. In the next two lemmas we state crucial properties for greedy decompositions of arbitrary (s,t)-flows.

Lemma 6 *Let $c \in \mathbb{R}^E$ and suppose that G is a series-parallel (s,t)-network. For an (s,t)-flow f of G of value N, let P be an arbitrary decomposition of f into N (s,t)-paths $\{p^1, \ldots, p^N\}$ with $c(p^i) \geq c(p^{i+1})$, $i \in [N-1]$, and let $\bar{P} = \{\bar{p}^1, \ldots, \bar{p}^N\}$ be a c-greedy decomposition of f. Then for all $x \geq 0$ we have that $R(P, c, x) \leq R(\bar{P}, c, x)$.*

Proof. The proof is by induction on the value N of the flow f. The base case is $N = 1$. In this case f is a single path, thus $P = \bar{P}$ and

$$R(P, c, x) \leq R(\bar{P}, c, x). \tag{2}$$

trivially holds.

Assume that (2) holds for $N \leq k$. When $N = k + 1$ we first prove that (2) holds in the case where $x = \bar{x}$, with $\bar{x} = \frac{c(f)}{N}$. We need the following claim.

Claim 1. *There is a decomposition* $\hat{P} = \{\hat{p}^1, \hat{p}^2, \ldots, \hat{p}^N\}$ *of* f *such that* $\hat{p}^1 = \bar{p}^1$ *and* $R(P, c, \bar{x}) \leq R(\hat{P}, c, \bar{x})$.

Proof of claim. For a decomposition P' of f into N (s, t)-paths let

$$\ell(P') = \min \left\{ |q \, \Delta \, \bar{p}^1| : q \in P', c(q) \geq \bar{x} \right\}.$$

Note that $\ell(P') = 0$ if and only if $\bar{p}^1 \in P'$. We want to prove that

$$\min\{\ell(P') : R(P, c, \bar{x}) \leq R(P', c, \bar{x}), \; P' \text{ decomposition of } f\} \tag{3}$$

is zero. Let $\tilde{P} = \{\tilde{p}^1, \ldots, \tilde{p}^N\}$ be a decomposition of f that achieves the minimum in (3) and assume by contradiction that $\ell(\tilde{P}) \geq 1$. Let π be an (s, t)-path of $\tilde{P}$ such that $c(\pi) \geq \bar{x}$ and $\ell(\tilde{P}) = |\pi \, \Delta \, \bar{p}^1|$. Since $\ell(\tilde{P}) \geq 1$, there exist two internally disjoint (u, v)-paths in $\pi \, \Delta \, \bar{p}^1$. We first restrict our attention to the set of paths $\tilde{P}_{uv}$ of $\tilde{P}$ that traverse nodes u and v, and to the corresponding (s, t)-flow $h = f(\tilde{P}_{uv})$. Note that $\pi \in \tilde{P}_{uv}$. Next, we will show how to construct a decomposition $\check{P}_{uv}$ of h such that $R(\check{P}_{uv}, c, \bar{x}) \geq R(\tilde{P}_{uv}, c, \bar{x})$ and $\ell(\check{P}_{uv}) < \ell(\tilde{P}_{uv})$. To this purpose, we will use an intermediate decomposition $Q \cup \pi$ of h, where $Q = \{q^1, \ldots, q^t\}$ is a c-greedy decomposition of $h \setminus f(\pi)$. Our target decomposition $\check{P}_{uv}$ will be obtained by slightly modifying this decomposition $Q \cup \pi$. We first prove that

$$R(Q \cup \pi, c, \bar{x}) \geq R(\tilde{P}_{uv}, c, \bar{x}).$$

In fact, since $t \leq k$, by our inductive hypothesis we have $R(\tilde{P}_{uv} \setminus \pi, c, \bar{x}) \leq R(Q, c, \bar{x})$, and by adding $c(\pi) - \bar{x}$ on both sides we obtain the above inequality.

We now specify how to construct $\check{P}_{uv}$. Let π_{uv} and q^1_{uv} be the (u, v)-subpaths of π and q^1, respectively. By construction, we have that q^1_{uv} is the (u, v)-subpath of $\bar{p}^1$. We define the decomposition $\check{P}_{uv}$ from $Q \cup \pi$ by replacing π with $\check{\pi} = \pi \setminus \pi_{uv} \cup q^1_{uv}$ and by replacing q^1 with $q^1 \setminus q^1_{uv} \cup \pi_{uv}$. This immediately implies

$$\ell(\check{P}_{uv}) \leq |\check{\pi} \, \Delta \, \bar{p}^1| < |\pi \, \Delta \, \bar{p}^1| = \ell(\tilde{P}_{uv}).$$

We prove that

$$R(\check{P}_{uv}, c, \bar{x}) \geq R(Q \cup \pi, c, \bar{x}). \tag{4}$$

Let $\delta = c(q^1_{uv}) - c(\pi_{uv})$. First, $\delta \geq 0$, because q^1_{uv} is a subpath of $\bar{p}^1$. Recalling that $c(\pi) \geq \bar{x}$, we obtain

$$R(\check{P}_{uv}, c, \bar{x}) = R(Q \cup \pi, c, \bar{x}) + \delta - \max\{0, c(q^1) - \bar{x}\} + \max\{0, c(q^1) - \delta - \bar{x}\}.$$

If $c(q^1) \leq \bar{x}$, since $\delta \geq 0$, we immediately have (4). If $c(q^1) \geq \bar{x}$ and $c(q^1) - \delta \geq \bar{x}$, we get

$$R(\check{P}_{uv}, c, \bar{x}) = R(Q \cup \pi, c, \bar{x}) + \delta - (c(q^1) - \bar{x}) + (c(q^1) - \delta - \bar{x})$$
$$= R(Q \cup \pi, c, \bar{x})$$

9

Finally, if $c(q^1) \geq \bar{x}$ and $c(q^1) - \delta \leq \bar{x}$, we get

$$
\begin{aligned}
R(\breve{P}_{uv}, c, \bar{x}) &= R(Q \cup \pi, c, \bar{x}) + \delta - (c(q^1) - \bar{x}) \\
&= R(Q \cup \pi, c, \bar{x}) - (c(q^1) - \delta - \bar{x}) \\
&\geq R(Q \cup \pi, c, \bar{x}).
\end{aligned}
$$

Thus, in any case (4) holds. We have obtained

$$
R(\breve{P}_{uv}, c, \bar{x}) \geq R(Q \cup \pi, c, \bar{x}) \geq R(\tilde{P}_{uv}, c, \bar{x}).
$$

We finally add to $\breve{P}_{uv}$ the paths in $\tilde{P} \setminus \tilde{P}_{uv}$. From the previous inequality we get

$$
R(\breve{P}, c, \bar{x}) \geq R(\tilde{P}, c, \bar{x}),
$$

where $\breve{P} = \breve{P}_{uv} \cup \tilde{P} \setminus \tilde{P}_{uv}$. Thus $R(\breve{P}, c, \bar{x}) \geq R(P, c, \bar{x})$. This is a contradiction on the choice of $\tilde{P}$, since

$$
\ell(\breve{P}) \leq \ell(\breve{P}_{uv}) < \ell(\tilde{P}_{uv}) = \ell(\tilde{P}).
$$

$\diamond$

With Claim 1 at hand we can prove that (2) holds for $N = k+1$ and $x = \bar{x}$. In fact, we only need to show that $R(\hat{P}, c, \bar{x}) \leq R(\bar{P}, c, \bar{x})$. We consider the (s,t)-flow $f \setminus \bar{p}^1$ of value k, and its decompositions $\hat{P} \setminus \bar{p}^1$ and $\bar{P} \setminus \bar{p}^1$. Note that $\bar{P} \setminus \bar{p}^1$ is a c-greedy decomposition of $f \setminus \bar{p}^1$. By induction we have that $R(\hat{P} \setminus \bar{p}^1, c, \bar{x}) \leq R(\bar{P} \setminus \bar{p}^1, c, \bar{x})$. By adding $c(\bar{p}^1) - \bar{x}$ on both sides we obtain $R(\hat{P}, c, \bar{x}) \leq R(\bar{P}, c, \bar{x})$, as desired.

We now prove that (2) holds for $N = k+1$ and $x \geq 0$. First, we remark that for each decomposition P' of f and $x \geq 0$ we have

$$
R(P', c, x) = \sum_{p \in P'} \max\{0, c(p) - x\} \geq \max\{0, c(f) - (k+1)x\}. \tag{5}
$$

Recall that the paths in P are listed in non-increasing order of cost. If $x \geq c(p^1)$, then $c(p^i) \leq c(p^1) \leq x$ for $i \in [N]$ implies $R(P, c, x) = 0$, and by (5) $R(\bar{P}, c, x) \geq 0$, thus (2) holds. If $x \leq c(p^{k+1})$, then $c(p^i) \geq c(p^{k+1}) \geq x$ for $i \in [N]$ implies $R(P, c, x) = c(f) - (k+1)x$, and by (5) $R(\bar{P}, c, x) \geq c(f) - (k+1)x$, thus (2) holds.

Thus we now assume $c(p^{k+1}) \leq x \leq c(p^1)$. Consider the network H obtained from G by adding $k+1$ parallel edges $e_1, \ldots, e_{k+1}$ from t to a new node t'. Define $c' \in \mathbb{R}^{E \cup \{e_1, \ldots, e_{k+1}\}}$ by setting $c'_e = c_e$ for $e \in E$, and

$$
c_{e_i} = \begin{cases} \max\{0, \alpha\} & i = 1 \\ \min\{0, \alpha\} & i = k+1 \\ 0 & i = 2, \ldots, k, \end{cases} \tag{6}
$$

where $\alpha = (k+1)x - c(f)$. Define the (s,t')-flow h of value $k+1$ obtained from f by assigning flow value 1 to all the new parallel edges $e_1, \ldots, e_{k+1}$. Finally, consider the decompositions $Q = \{q^1, \ldots, q^{k+1}\}$ and $\bar{Q} = \{\bar{q}^1, \ldots, \bar{q}^{k+1}\}$ obtained from P and $\bar{P}$, respectively, by appending edge e_i to the i-th paths of the decompositions. More precisely,

$q^i = p^i \circ e_i$ and $\bar{q}^i = \bar{p}^i \circ e_i$ for $i \in [k+1]$. First, we remark that $x = \frac{c'(h)}{k+1}$. Secondly, by construction $\bar{Q}$ is a c'-greedy decomposition of h. Since we have proven that (2) holds for $N = k+1$ and $x = \frac{c'(h)}{N}$, we have

$$R(Q, c', x) \leq R(\bar{Q}, c', x).$$

If $x \geq \frac{c(f)}{k+1}$ we have $\alpha \geq 0$, thus e_1 has nonnegative cost and $e_2, \ldots, e_{k+1}$ have costs 0. Thus

$$
\begin{aligned}
c'(q^1) &= c(p^1) + \alpha \\
c'(\bar{q}^1) &= c(\bar{p}^1) + \alpha \\
c'(q^i) &= c(p^i) & i &= 2, \ldots, k+1 \\
c'(\bar{q}^i) &= c(\bar{p}^i) & i &= 2, \ldots, k+1.
\end{aligned}
$$

Since, by Lemma 5, $x \leq c(p^1) \leq c(\bar{p}^1)$ we have

$$
\begin{aligned}
R(Q, c', x) &= R(P, c, x) + \alpha \\
R(\bar{Q}, c', x) &= R(\bar{P}, c, x) + \alpha,
\end{aligned}
$$

thus (2) holds.

If $\max\{0, c(p^{k+1})\} \leq x \leq \frac{c(f)}{k+1}$ we have $\alpha \leq 0$, thus e_{k+1} has nonpositive cost and $e_1, \ldots, e_k$ have costs 0. Thus

$$
\begin{aligned}
c'(q^i) &= c(p^i) & i &\in [k] \\
c'(\bar{q}^i) &= c(\bar{p}^i) & i &\in [k] \\
c'(q^{k+1}) &= c(p^{k+1}) + \alpha \\
c'(\bar{q}^{k+1}) &= c(\bar{p}^{k+1}) + \alpha
\end{aligned}
$$

Since, by Lemma 5, $x \geq c(p^{k+1}) \geq c(\bar{p}^{k+1})$ we have

$$
\begin{aligned}
R(Q, c', x) &= R(P, c, x) \\
R(\bar{Q}, c', x) &= R(\bar{P}, c, x),
\end{aligned}
$$

thus (2) holds.

$\square$

Lemma 7 *Suppose that G is a series-parallel (s, t)-network and that the delay functions are affine. Let $c_e = d_e(f_e)$ for all $e \in E$. Let f be a PNE and $\bar{P} = \bar{P}(f) = \{\bar{p}^1, \bar{p}^2, \ldots, \bar{p}^N\}$ be a c-greedy decomposition of f. Then $c(\bar{p}^{i+1}) \geq \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}^j)}{i}$ for $i \in [N-1]$.*

Proof. We prove the lemma by induction on the number of edges $|E|$ in G. In the base case, we have $|E| = 1$, i.e., G is a single edge, thus for every $\bar{p}^i \in \bar{P}$ we have $c(\bar{p}^1) = c(\bar{p}^2) = \cdots = c(\bar{p}^N)$ and we are done. Next we assume that when $|E| \leq k$, any c-greedy decomposition $\bar{P}$ of the PNE flow f on G satisfies Lemma 7. We need to prove

11

that when $|E| = k + 1$, Lemma 7 still holds. Because G is a series-parallel network and $|E| > 1$, G must be composed in series or in parallel by two series-parallel subgraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$. We have $|E_1| < |E|$ and $|E_2| < |E|$. We also denote the subflow of f in G_1 by f_1, the subflow of f in G_2 by f_2.

We denote by $P^* = \{p^1, p^2, \dots, p^N\}$ a decomposition of f defining the single players' strategies of a PNE. We first outline the proof structure. The main idea is to break down P^* and $\bar{P}$ according to the decomposition of G into G_1, G_2. We will first prove that f_1 and f_2 are PNE in G_1 and G_2, respectively. Then we will argue that the decompositions $\bar{P}_1$ and $\bar{P}_2$ of f_1 and f_2 obtained by breaking down $\bar{P}$ are c-greedy decompositions in G_1, G_2, respectively. Finally, we will apply induction. We consider separately the cases in which G is composed in series and in parallel.

Case 1 : G is composed by G_1, G_2 in series at node a. Then also $G(f)$ is obtained by composing in series $G(f_1), G(f_2)$ at node a.

We first show that f_1 is a PNE flow in G_1 and f_2 is a PNE flow in G_2. For every (s,t)-path $p^i \in P^*$ we have $p^i = p_1^i \circ p_2^i$, where p_1^i is an (s,a)-path in G_1 and p_2^i is an (a,t)-path in G_2. Thus $P_1^* = \{p_1^1, p_1^2, \dots, p_1^N\}$ is a decomposition of f_1, $P_2^* = \{p_2^1, p_2^2, \dots, p_2^N\}$ is a decomposition of f_2. If f_1 is not a PNE flow in G_1, there exists a $p_1^i \in P_1^*$ such that $c(p_1^i)$ will decrease if we switch p_1^i to some (s,a)-path q on G_1. Now consider the (s,t)-path $p^i \in P^*$, and note that $c(p^i)$ will also decrease if we switch to the (s,t)-path $q \circ p_2^i$ on G. This contradicts the fact that P^* is a PNE in G. The proof for f_2 is similar.

Moreover, for each $\bar{p}^i \in \bar{P}$, we have that $\bar{p}^i = \bar{p}_1^i \circ \bar{p}_2^i$, where $\bar{p}_1^i$ is an (s,a)-path and $\bar{p}_2^i$ is an (a,t)-path. We can conclude that $\bar{P}_1 = \{\bar{p}_1^1, \bar{p}_1^2, \dots, \bar{p}_1^N\}$ and $\bar{P}_2 = \{\bar{p}_2^1, \bar{p}_2^2, \dots, \bar{p}_2^N\}$ are c-greedy decompositions of f_1 and f_2 respectively, otherwise $\bar{P}$ would not be a c-greedy decomposition.

Then by our inductive hypothesis, we know that $\bar{P}_1$ and $\bar{P}_2$ satisfy Lemma 7 because $|E_1| < |E| = k + 1$ and $|E_2| < |E| = k + 1$. Thus $c(\bar{p}_1^{i+1}) \geq \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}_1^j)}{i}$ and $c(\bar{p}_2^{i+1}) \geq \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}_2^j)}{i}$ for $i \in [N-1]$. Note that because $\bar{p}^{i+1} = \bar{p}_1^{i+1} \circ \bar{p}_2^{i+1}$, we have that $c(\bar{p}^{i+1}) = c(\bar{p}_1^{i+1}) + c(\bar{p}_2^{i+1})$. Thus

$$c(\bar{p}^{i+1}) = c(\bar{p}_1^{i+1}) + c(\bar{p}_2^{i+1})$$
$$\geq \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}_1^j)}{i} + \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}_2^j)}{i} = \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}^j)}{i}.$$

Case 2 : G is composed by G_1, G_2 in parallel. Then also $G(f)$ is obtained by composing in parallel $G(f_1), G(f_2)$.

We first show that f_1 is a PNE flow in G_1 and f_2 is a PNE flow in G_2. Define $P_1^* = \{p^i \in P^* : p^i \text{ is a path in } G_1\}$ and $P_2^* = \{p^i \in P^* : p^i \text{ is a path in } G_2\}$. Note that $P^* = P_1^* \cup P_2^*$, P_1^* is a decomposition of f_1, and P_2^* is a decomposition of f_2. If f_1 is not a PNE flow in G_1, there exists a $p_1^i \in P_1^*$ such that $c(p_1^i)$ will decrease if we switch p_1^i to some (s,t)-path q in G_1. Since $p_1^i \in P^*$ and q is an (s,t)-path in G, this contradicts the fact that f is a PNE flow in G. The proof for f_2 is similar.

Moreover, from the c-greedy decomposition $\bar{P}$ of f we define

$$\bar{P}_1 = \{p \in \bar{P}(f) : p \text{ is a path in } G_1\}$$
$$\bar{P}_2 = \{p \in \bar{P}(f) : p \text{ is a path in } G_2\},$$

that are c-greedy decompositions of f_1 and f_2, respectively. For $i \in [N-1]$ define

$$\bar{P}^i = \{\bar{p}^j \in \bar{P}(f) : j \leq i\}$$
$$\bar{P}_1^i = \{\bar{p}^j \in \bar{P}^i : p \text{ is a path in } G_1\}$$
$$\bar{P}_2^i = \{\bar{p}^j \in \bar{P}^i : p \text{ is a path in } G_2\}.$$

Note that $\bar{P}^i = \bar{P}_1^i \cup \bar{P}_2^i$. Our goal is to prove

$$c(\bar{p}^{i+1}) \geq \frac{1}{2} \sum_{j=1}^{i} \frac{c(\bar{p}^j)}{i}$$
$$= \frac{1}{2} \frac{\sum_{p \in \bar{P}_1^i} c(p) + \sum_{p \in \bar{P}_2^i} c(p)}{|\bar{P}_1^i| + |\bar{P}_2^i|}. \tag{7}$$

Assume without loss of generality that $\bar{p}^{i+1} \in \bar{P}_1$. Since f_1 is a PNE flow in G_1, $\bar{P}_1$ is a c-greedy decomposition of f_1, and $|E_1| < |E| = k+1$, by our inductive hypothesis we have

$$c(\bar{p}^{i+1}) \geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_1^i} c(p)}{|\bar{P}_1^i|}. \tag{8}$$

If $\bar{P}_2^i = \emptyset$, (7) trivially holds. Thus we now assume $\bar{P}_2^i \neq \emptyset$. In this setting, there are some paths of $\bar{P}^i$ that belong to G_2. Thus f_1 and f_2 are nonzero flows, we must have $f_e \leq N-1$ for all $e \in E$. In the rest of the proof we show

$$c(\bar{p}^{i+1}) \geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_2^i} c(p)}{|\bar{P}_2^i|}. \tag{9}$$

Note that (8) and (9) imply (7).

If $|\bar{P}_2^i| = |\bar{P}_2|$, all the paths of $\bar{P}(f)$ that belong to G_2 appear before $\bar{p}^{i+1}$. So it is sufficient to show that $c(\bar{p}^{i+1}) \geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_2} c(p)}{|\bar{P}_2|}$. We have:

$$c(\bar{p}^{i+1}) = \sum_{e \in \bar{p}^{i+1}} d_e(f_e)$$
$$\geq \frac{1}{2} \sum_{e \in \bar{p}^{i+1}} d_e(f_e + 1) \tag{10}$$
$$\geq \frac{1}{2} \max \{c(p) : p \in P_2^*\} \tag{11}$$
$$\geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_2} c(p)}{|\bar{P}_2|}. \tag{12}$$

13

Inequality (10) holds since for each edge $e \in E$ we have an affine delay function $d_e(x) = a_e x + b_e$ with $a_e \geq 0, b_e \geq 0$. More precisely, when $a_e = b_e = 0$ we have $d_e(f_e) = d_e(f_e+1) = 0$, and when at least one among a_e and b_e is positive we have $\frac{d_e(f_e)}{d_e(f_e+1)} = \frac{a_e f_e + b_e}{a_e(f_e+1)+b_e} \geq \frac{1}{2}$ for any $f_e \in [N-1]$. Inequality (11) holds since P^* is a PNE and because $\bar{p}^{i+1}$ is an (s,t)-path. Finally, inequality (12) holds since the cost of the most expensive path in P_2^* is higher that the average cost of the paths in P_2^*, which is equal to the average cost of the paths in $\bar{P}_2$.

If $|\bar{P}_2^i| < |\bar{P}_2|$, some of the paths of $\bar{P}$ that belong to G_2 appear before $\bar{p}^{i+1}$ and some appear after $\bar{p}^{i+1}$. We denote by $\ell = \min\{t : \bar{p}^t \in \bar{P}_2, t > i+1\}$. Then we have $c(\bar{p}^{i+1}) \geq c(\bar{p}^\ell)$ because $\ell > i+1$. Since $|E_2| < |E| = k+1$, by our inductive hypothesis, we have

$$c(\bar{p}^{i+1}) \geq c(\bar{p}^\ell) \geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_2^{\ell-1}} c(p)}{|\bar{P}_2^{\ell-1}|} \geq \frac{1}{2} \frac{\sum_{p \in \bar{P}_2^i} c(p)}{|\bar{P}_2^i|},$$

which proves (9). This completes the proof. $\qquad\square$

Example 1, continued *We illustrate the properties stated in Lemmas 6 and 7 on the example in Fig. 1. First, the average cost $\bar{x} = c(f)/N = 5$. We have $R(P, c, \bar{x}) = (5-5) + (5-5) + (5-5) = 0$ and $R(\bar{P}, c, \bar{x}) = (6-5) + (6-5) + 0 = 2$. First, Lemma 6 holds since $R(P, c, \bar{x}) = 0 < 2 = R(\bar{P}, c, \bar{x})$. For Lemma 7, we observe that the paths in the c-greedy decomposition $\bar{P} = \{\bar{p}^1, \bar{p}^2, \bar{p}^3\}$ have costs $6, 6$ and 3, respectively, thus $c(\bar{p}^2) = 6 \geq 3 = \frac{1}{2} c(\bar{p}^1)$ and $c(\bar{p}^3) = 3 \geq 3 = \frac{1}{2} \frac{c(\bar{p}^1) + c(\bar{p}^2)}{2}$.*

Based on Lemma 7, from $\bar{P}(f) = \{\bar{p}^1, \bar{p}^2, \ldots, \bar{p}^N\}$, we get a sequence of positive numbers $\{\text{cost}_f(\bar{p}^1), \text{cost}_f(\bar{p}^2), \ldots, \text{cost}_f(\bar{p}^N)\}$ such that $\text{cost}_f(\bar{p}^{i+1}) \geq \frac{1}{2} \sum_{j=1}^i \frac{\text{cost}_f(\bar{p}^j)}{i}$, $i \in [N-1]$. We now turn our attention to general sequences of positive numbers that satisfy this property. For $m \in [N-1]$ we define $\mu(m, N) = \prod_{j=m}^{N-1} \frac{2j}{2j+1}$.

Lemma 8 *Let $x \in \mathbb{R}_+^N$ such that $\sum_{i=1}^N x_i = 1$, and let $m \in [N-1]$. We have:*

1. *If $x_{i+1} \geq \frac{1}{2i} \sum_{j=1}^i x_j$ for $i \in [N-1]$, then $\sum_{i=1}^m x_i \leq \mu(m, N)$.*

2. *If $x_{i+1} = \frac{1}{2i} \sum_{j=1}^i x_j$ for $i \in [N-1]$, then $\sum_{i=1}^m x_i = \mu(m, N)$.*

Proof.

We first prove statement 1. We proceed by induction on m. The base case is $m = N-1$. Since $x_N \geq \frac{1}{2(N-1)} \sum_{j=1}^{N-1} x_j$, we have:

$$\sum_{j=1}^{N-1} x_j = 1 - x_N \leq 1 - \frac{1}{2(N-1)} \sum_{j=1}^{N-1} x_j. \qquad (13)$$

By equation (13), we have $\frac{2(N-1)+1}{2(N-1)} \sum_{j=1}^{N-1} x_j \leq 1$. This implies that $\sum_{j=1}^{N-1} x_j \leq \frac{2(N-1)}{2(N-1)+1} = \mu(N-1, N)$. Thus statement 1 holds for the base case.

Next we assume that statement 2 holds for $m \in \{k, \ldots, N-1\}$, and we prove that it also holds for $m = k - 1$. Based on our inductive hypothesis, $\sum_{j=1}^{k} x_j \leq \mu(k, N)$. Moreover, since $x_k \geq \frac{1}{2(k-1)} \sum_{j=1}^{k-1} x_j$, we have:

$$\sum_{j=1}^{k-1} x_j = \sum_{j=1}^{k} x_j - x_k \leq \mu(k, N) - \frac{1}{2(k-1)} \sum_{j=1}^{k-1} x_j. \tag{14}$$

According to (14), we have $\frac{2(k-1)+1}{2(k-1)} \sum_{j=1}^{k-1} x_j \leq \mu(k, N)$. This implies that $\sum_{j=1}^{k-1} x_j \leq \frac{2(k-1)}{2(k-1)+1} \mu(k, N) = \mu(k-1, N)$. Thus, statement 1 holds.

The proof of statement 2 is analogous, and it is obtained by replacing the inequalities in (13) and (14) with equalities. $\qquad \square$

The next lemma provides a lower and an upper bound for $\mu(m, N)$.

Lemma 9 *For $m \in [N-1]$ we have*

$$\sqrt{\frac{2m-1}{2N-1}} \leq \mu(m, N) \leq \sqrt{\frac{m}{N}}.$$

Proof. First we can equivalently write:

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

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

$$\prod_{j=m}^{N-1} \left(\frac{2j}{2j+1} \right)^2 \geq \frac{2m-1}{2m} \cdot \frac{2m}{2m+1} \cdot \frac{2m+1}{2m+2} \cdot \frac{2m+2}{2m+3} \cdots \frac{2N-3}{2N-2} \cdot \frac{2N-2}{2N-1}$$

$$= \frac{2m-1}{2N-1}.$$

Similarly, we upper bound the argument of the square root as follows.

$$\prod_{j=m}^{N-1} \left(\frac{2j}{2j+1} \right)^2 \leq \frac{2m}{2m+1} \cdot \frac{2m+1}{2m+2} \cdot \frac{2m+2}{2m+3} \cdot \frac{2m+3}{2m+4} \cdots \frac{2N-2}{2N-1} \cdot \frac{2N-1}{2N}$$

$$= \frac{2m}{2N}.$$

$\qquad \square$

From the previous results we can establish the following property of a PNE f, which will be used to prove Theorem 3.

Lemma 10 *Suppose that G is a series-parallel (s,t)-network and that the delay functions are affine. Let $c_e = d_e(f_e)$ for all $e \in E$. Moreover, let f be a PNE and let $\bar{P} = \{\bar{p}^1, \ldots, \bar{p}^N\}$ be a c-greedy decomposition of f. We have*

$$R\left(\bar{P}, c, \frac{cost(f)}{N}\right) \leq \frac{1}{4} cost(f).$$

Proof. Let m be the number of paths in $\bar{P}$ whose cost is greater than $\text{cost}(f)/N$ and note that

$$R\left(\bar{P}, c, \frac{\text{cost}(f)}{N}\right) = \sum_{i=1}^{m} c(\bar{p}^i) - \frac{m}{N}\text{cost}(f).$$

We equivalently prove

$$\sum_{i=1}^{m} c(\bar{p}^i) \leq \left(\frac{1}{4} + \frac{m}{N}\right) \text{cost}(f). \tag{15}$$

Define $x_i = c(p^i)/\text{cost}(f)$ for $i \in [N]$. Clearly $x_i \geq 0$ for $i \in [N]$ and $\sum_{i=1}^{N} x_i = 1$. By Lemma 7 we know that $c(\bar{p}^{i+1}) \geq \frac{1}{2}\sum_{j=1}^{i} \frac{c(\bar{p}^j)}{i}$ for $i \in [N-1]$, thus $x_{i+1} \geq \frac{1}{2}\sum_{j=1}^{i} \frac{x_j}{i}$ for $i \in [N-1]$. By Lemmas 8 and 9 we have

$$\sum_{i=1}^{m} c(p^i) = \text{cost}(f)\sum_{i=1}^{m} x_i \leq \mu(m,N)\text{cost}(f) \leq \sqrt{\frac{m}{N}}\text{cost}(f).$$

To show (15), we finally observe that $\sqrt{\frac{m}{N}} - \frac{m}{N} \leq \frac{1}{4}$, since $0 \leq \frac{m}{N} \leq 1$ and $\sqrt{x} - x \leq \frac{1}{4}$ for $x \in [0,1]$. $\qquad\square$

Finally, we state the following elementary property of a PNE flow in a series-parallel congestion game.

Lemma 11 *Suppose that G is a series-parallel (s,t)-network. Let f be a PNE flow of value N and let p be an (s,t)-path in G. Then $cost_f^+(p) \geq \frac{cost(f)}{N}$.*

Proof. Denote by P^* the set of N (s,t)-paths in the PNE associated to f. Clearly $\max\{\text{cost}_f(\pi) : \pi \in P^*\} \geq \frac{\text{cost}(f)}{N}$. By contradiction, suppose that $\text{cost}_f^+(p) < \frac{\text{cost}(f)}{N}$. We would obtain that $\max\{\text{cost}_f(\pi) : \pi \in P^*\} > \text{cost}_f^+(p)$, thus one player would find profitable to change her strategy into p. This contradicts the fact that f is a PNE. $\qquad\square$

We are finally ready to prove the central result of our paper.

Proof of Theorem 3. If $\mathcal{C} = \emptyset$, the claim trivially holds since $\Delta(\mathcal{C}, f) = 0$. Thus, we now assume $\mathcal{C} \neq \emptyset$. Let $\mathcal{C}^- = \{C_1^-, \ldots, C_k^-\}$. The proof is by induction on

$$\gamma(\mathcal{C}^-, G) = \min\left\{\sum_{i=1}^{k} |p_i \setminus C_i^-| : p_i \text{ is an } (s,t)\text{-path containing } C_i^-, i \in [k]\right\}. \tag{16}$$

The base case is $\gamma(\mathcal{C}^-, G) = 0$, in which case for all $i \in [k]$ we have $p_i = C_i^-$, i.e., C_i^- is an (s,t)-path. Let P be a decomposition of f containing all the paths in $\mathcal{C}^-$, and let $\bar{P}$ be a c-greedy decomposition of f, where $c_e = d_e(f_e)$ for all $e \in E$. We obtain:

$$\Delta(\mathcal{C}, f) = \sum_{i=1}^{k} (\mathrm{cost}_f(C_i^-) - \mathrm{cost}_f^+(C_i^+))$$

$$\leq \sum_{i=1}^{k} \left(\mathrm{cost}_f(C_i^-) - \frac{\mathrm{cost}(f)}{N} \right) \tag{17}$$

$$\leq R\left(\mathcal{C}^-, c, \frac{\mathrm{cost}(f)}{N} \right) \tag{18}$$

$$\leq R\left(P, c, \frac{\mathrm{cost}(f)}{N} \right) \tag{19}$$

$$\leq R\left(\bar{P}, c, \frac{\mathrm{cost}(f)}{N} \right) \tag{20}$$

$$\leq \frac{1}{4} \mathrm{cost}(f). \tag{21}$$

Inequality (17) holds since for all $i \in [k]$ C_i^+ is an (s,t)-path whose only nodes in common with C_i^- are s and t. Thus, by Lemma 11, we have $\mathrm{cost}_f^+(C_i^+) \geq \frac{\mathrm{cost}(f)}{N}$ for all $i \in [k]$. Inequality (18) follows from the definition of c and the fact that $R\left(\mathcal{C}^-, c, \frac{\mathrm{cost}(f)}{N}\right)$ only contains the nonnegative terms of the summation in (17). Inequality (19) holds since $\mathcal{C}^- \subseteq P$. Inequality (20) holds because of Lemma 6. Inequality (21) is implied by Lemma 10.

Now we assume that our claim holds if $\gamma(\mathcal{C}^-, G) \leq \bar{\gamma}$. Our goal is to show that the claim still holds if $\gamma(\mathcal{C}^-, G) = \bar{\gamma} + 1$. To prove that

$$\frac{\Delta(\mathcal{C}, f)}{\mathrm{cost}(f)} \leq \frac{1}{4}$$

we construct another instance of a N-player network congestion game where

(i) $\hat{G}$ is a series-parallel (s,t)-network with affine delays,

(ii) $\hat{f}$ is an (s,t)-flow of value N in $\hat{G}$ and $\hat{P}$ is a decomposition of $\hat{f}$ that is a PNE,

(iii) $\hat{\mathcal{C}} = \{(\hat{C}_i^-, \hat{C}_i^+) : i \in [h]\}$ is a collection of h pairs of internally disjoint paths in $\hat{G}$ with $\left| \left\{ \hat{C}_i^- : e \in \hat{C}_i^- \right\} \right| \leq \hat{f}_e$ for all $e \in E(\hat{G})$,

(iv) $\gamma(\hat{\mathcal{C}}^-, \hat{G}) \leq \bar{\gamma}$, where $\hat{\mathcal{C}}^- = \{\hat{C}_1^-, \ldots, \hat{C}_h^-\}$,

(v) $\dfrac{\Delta(\mathcal{C}, f)}{\mathrm{cost}(f)} \leq \dfrac{\Delta(\hat{\mathcal{C}}, \hat{f})}{\mathrm{cost}(\hat{f})}$.
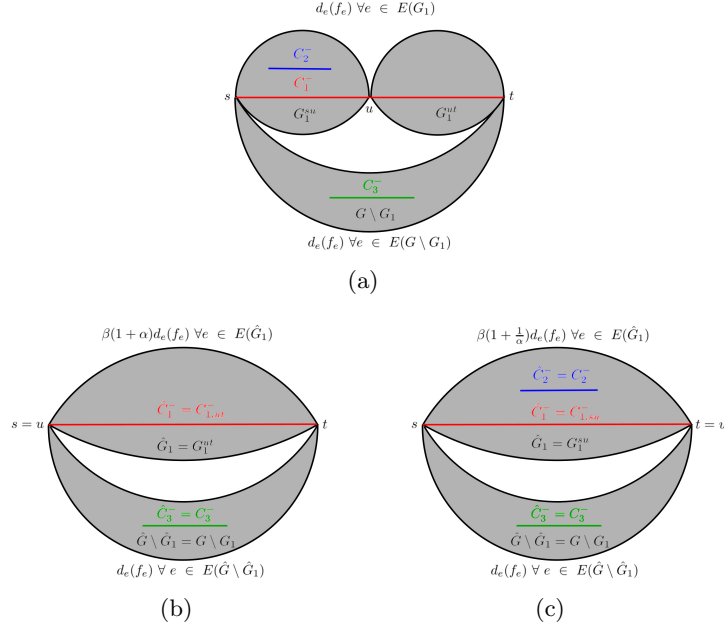
Figure 2: This is an example for the operations in Theorem 3, where the dotted lines are paths C_i^- in the set $\mathcal{C}^-$. (a) A series-parallel network G. (b) The network $\hat{G}$ by applying Operation 1 to G and (c) the network $\hat{G}$ by applying Operation 2 to G.

Intuitively, by decreasing $\gamma(\mathcal{C}^-, G)$ at each step we reduce, in a finite number of steps, to a network $\hat{G}$ where the number of non-(s,t)-paths in $\mathcal{C}^-$ has strictly decreased. First, we describe how to construct $\hat{G}$, $\hat{f}$, $\hat{P}$ and $\hat{\mathcal{C}}$.

Let G be composed in parallel by $G_1, \ldots, G_\ell$, $\ell \geq 1$, and assume wlog that each G_i cannot be further decomposed in parallel. Since $\gamma(\mathcal{C}^-, G) \geq 1$, there is at least a (w,v)-path C_j^- in $\mathcal{C}^-$ that is not from s to t. We assume wlog that C_j^- is contained in G_1, and we define f_1 to be the subflow of f in G_1. Since C_j^- is not from s to t, G_1 can be decomposed in series. Moreover, since C_j^- and C_j^+ are internally disjoint, there must be a component of the series decomposition of G_1 which contains C_j^-. Thus there exists a node $u \in V(G_1)$ such that G_1 is obtained by composing in series at u two subgraphs G_1^{su} and G_1^{ut}, and C_j^- is contained either in G_1^{su} or in G_1^{ut}. Correspondingly, we can also split the flow f_1 into an (s,u)-flow f_1^{su} and a (u,t)-flow f_1^{ut}.

Let $\mathcal{C}^-(G_1)$ consist of the paths in $\mathcal{C}^-$ that are contained in G_1. Analogously, let $\mathcal{C}^-(G_1^{su})$ and $\mathcal{C}^-(G_1^{ut})$ be the paths of $\mathcal{C}^-(G_1)$ that are contained in G_1^{su} and G_1^{ut}, respectively. Note that each path C_j^- that is contained in $\mathcal{C}^-(G_1) \setminus (\mathcal{C}^-(G_1^{su}) \cup \mathcal{C}^-(G_1^{ut}))$ must be an (s,t)-path, since otherwise the path C_j^+ would also belong to G_1 and thus traverse u, contradicting the assumption that C_j^- and C_j^+ are internally disjoint.

Let $P^* = \{p^1, \ldots, p^N\}$ be the (s,t)-paths chosen by the players in the PNE f. Let $\alpha > 0$, $\beta > 0$. We define two operations, whose pictorial representations are given in Fig. 2.

18

Operation 1.

1. Define network $\hat{G}$ obtained from G by shrinking G_1^{su} (G_1 is replaced by G_1^{ut}, and nodes s and u are identified).

2. For each edge e of $\hat{G}$ that is in G_1^{ut}, redefine the delay to be $\beta(1 + \alpha)d_e(f_e)$.

3. Construct an (s,t)-flow $\hat{f}$ of value N in $\hat{G}$ from f, by replacing f_1 with f_1^{ut}. Set $\hat{P} = \{\hat{p}^1, \ldots, \hat{p}^N\}$ to be the (s,t)-paths chosen by the players in $\hat{f}$, where $\hat{p}^i$ is the (u,t)-subpath of p^i if p^i is in G_1 and $\hat{p}^i = p^i$ if p^i is not in G_1.

4. Define set $\hat{\mathcal{C}}$ containing:

   (a) all $(C^-, C^+) \in \mathcal{C}$ such that $C^- \notin \mathcal{C}^-(G_1)$.

   (b) all $(C^-, C^+) \in \mathcal{C}$ such that $C^- \in \mathcal{C}^-(G_1^{ut})$.

   (c) all (C_{ut}^-, C^+), such that $(C^-, C^+) \in \mathcal{C}$, C^- is an (s,t)-path in G_1, and C_{ut}^- is the subpath of C^- from u to t.

Operation 2.

1. Define network $\hat{G}$ obtained from G by shrinking G_1^{ut} (G_1 is replaced by G_1^{su}, and nodes u and t are identified).

2. For each edge e of $\hat{G}$ that is in G_1^{su}, redefine the delay to be $\beta(1 + \frac{1}{\alpha})d_e(f_e)$.

3. Construct an (s,t)-flow $\hat{f}$ of value N in $\hat{G}$ from f, by replacing f_1 with f_1^{su}. Set $\hat{P} = \{\hat{p}^1, \ldots, \hat{p}^N\}$ to be the (s,t)-paths chosen by the players in $\hat{f}$, where $\hat{p}^i$ is the (s,u)-subpath of p^i if p^i is in G_1 and $\hat{p}^i = p^i$ if p^i is not in G_1.

4. Define set $\hat{\mathcal{C}}$ containing:

   (a) all $(C^-, C^+) \in \mathcal{C}$ such that $C^- \notin \mathcal{C}^-(G_1)$.

   (b) all $(C^-, C^+) \in \mathcal{C}$ such that $C^- \in \mathcal{C}^-(G_1^{su})$.

   (c) all (C_{su}^-, C^+), such that $(C^-, C^+) \in \mathcal{C}$, C^- is an (s,t)-path in G_1, and C_{su}^- is the subpath of C^- from s to u.

The network $\hat{G}$ obtained with Operation 1 (resp. Operation 2) is a series-parallel (s,t)-network with affine delays, thus (i) is satisfied. For $i \in [k]$, denote by $\mathrm{cost}_f(C_i)$ the difference $\mathrm{cost}_f(C_i^-) - \mathrm{cost}_f^+(C_i^+)$. Let

$$D = \sum_{C_i^- \in \mathcal{C}^-(G_1^{su})} \mathrm{cost}_f(C_i) + \sum_{C_i^- \in \mathcal{C}^-(G_1^{ut})} \mathrm{cost}_f(C_i) + \sum_{C_i^- \in \mathcal{C}^-(G_1) \ (s,t)\text{-path}} \mathrm{cost}_f(C_i^-).$$

The next claim shows that (v) is also satisfied by appropriately performing either Operation 1 or Operation 2.

Claim 2. *If*

$$\frac{\Delta(\mathcal{C}, f)}{cost(f)} \leq \frac{D}{cost(f_1)}, \tag{22}$$

then for each $\beta \geq 1$ and $\alpha > 0$ either Operation 1 or Operation 2 yields

$$\frac{\Delta(\mathcal{C}, f)}{cost(f)} \leq \frac{\Delta(\hat{\mathcal{C}}, \hat{f})}{cost(\hat{f})}. \tag{23}$$

Otherwise, if inequality (22) does not hold, for each $\beta \leq 1$ and $\alpha > 0$ either Operation 1 or Operation 2 yields (23).

Proof of claim. If inequality (22) holds and we choose $\beta \geq 1$, then we have:

$$\frac{\Delta(\mathcal{C}, f)}{cost(f)} \leq \frac{\Delta(\mathcal{C}, f) + (\beta - 1)D}{cost(f) + (\beta - 1)cost(f_1)}. \tag{24}$$

If inequality (22) does not hold and we choose $\beta \leq 1$, then (24) still holds. Define

$$B(\xi) = \xi cost(f_1^{ut}) - cost(f_1^{su})$$

$$A(\xi) = - \sum_{C_i^- \in \mathcal{C}^-(G_1^{su})} cost_f(C_i) + \xi \sum_{C_i^- \in \mathcal{C}^-(G_1^{ut})} cost_f(C_i)$$

$$- \sum_{C_i^- \in \mathcal{C}^-(G_1)(s,t)\text{-path}} \left(cost_f((C_i^-)_{su}) - \xi cost_f((C_i^-)_{ut}) \right).$$

It can be checked that if we apply Operation 1 with parameters (α, β) we get:

$$\Delta(\hat{\mathcal{C}}, \hat{f}) = \Delta(\mathcal{C}, f) + A(\beta(1 + \alpha) - 1) = \Delta(\mathcal{C}, f) + (\beta - 1)D + \beta A(\alpha)$$

$$cost(\hat{f}) = cost(f) + B(\beta(1 + \alpha) - 1) = cost(f) + (\beta - 1)cost(f_1) + \beta B(\alpha).$$

Moreover, if we apply Operation 2 with parameters (α, β) we get:

$$\Delta(\hat{\mathcal{C}}, \hat{f}) = \Delta(\mathcal{C}, f) + A\left(\beta - 1 + \frac{\beta}{\alpha}\right) = \Delta(\mathcal{C}, f) + (\beta - 1)D - \frac{\beta}{\alpha}A(\alpha)$$

$$cost(\hat{f}) = cost(f) + B\left(\beta - 1 + \frac{\beta}{\alpha}\right) = cost(f) + (\beta - 1)cost(f_1) - \frac{\beta}{\alpha}B(\alpha).$$

Thus, if

$$\frac{A(\alpha)}{B(\alpha)} \geq \frac{\Delta(\mathcal{C}, f) + (\beta - 1)D}{cost(f) + (\beta - 1)cost(f_1)}$$

by applying Operation 1 with parameters (α, β) we get:

$$\frac{\Delta(\mathcal{C}, f) + (\beta - 1)D}{cost(f) + (\beta - 1)cost(f_1)} \leq \frac{\Delta(\mathcal{C}, f) + (\beta - 1)D + \beta A(\alpha)}{cost(f) + (\beta - 1)cost(f_1) + \beta B(\alpha)} = \frac{\Delta(\hat{\mathcal{C}}, \hat{f})}{cost(\hat{f})}.$$

Otherwise, by applying Operation 2 with parameters (α, β) we get:

$$\frac{\Delta(\mathcal{C}, f) + (\beta - 1)D}{cost(f) + (\beta - 1)cost(f_1)} \leq \frac{\Delta(\mathcal{C}, f) + (\beta - 1)D - \frac{\beta}{\alpha}A(\alpha)}{cost(f) + (\beta - 1)cost(f_1) - \frac{\beta}{\alpha}B(\alpha)} = \frac{\Delta(\hat{\mathcal{C}}, \hat{f})}{cost(\hat{f})}.$$

By choosing β appropriately, by (24) we have the desired result. $\diamond$

In the next two claims we show that if we apply Operation 1 (resp. 2) with appropriate parameters α and β, then also (ii) is satisfied. Let H be a subgraph of G that is a two-terminal series-parallel network with terminals u and v, and let P be a set of (u, v)-paths in H. We define

$$c(H) = \min\{\mathrm{cost}_f^+(p) : p \text{ is an } (u, v)\text{-path in } H\},$$
$$C(P) = \max\{\mathrm{cost}_f(p) : p \in P\}.$$

Let P_1^* and $\hat{P}_1$ be the paths in P^* and $\hat{P}$ that are contained in G_1 and $\hat{G}_1$, respectively. We denote by $P_{1,su}^*$ and $P_{1,ut}^*$ the set of (s, u)-subpaths of the paths in P_1^* and the set of (u, t)-subpaths of the paths in P_1^*, respectively. We define

$$\alpha^{\min} = \frac{c(G_1^{su})}{c(G_1^{ut})}, \quad \beta^{\min} = \frac{c(G)}{c(G_1)}, \quad \alpha^{\max} = \frac{C(P_{1,su}^*)}{C(P_{1,ut}^*)}, \quad \beta^{\max} = \frac{C(P^*)}{C(P_1^*)}.$$

Claim 3. *The decomposition of $\hat{f}$ into $\{\hat{p}^1, \ldots, \hat{p}^N\}$ obtained by applying Operation 1 (resp. Operation 2) with $(\alpha, \beta) = (\alpha^{min}, \beta^{min})$ is a PNE in the network congestion game on $\hat{G}$.*

Proof of claim. Suppose we apply Operation 1 (resp. Operation 2) with $(\alpha, \beta) = (\alpha^{\min}, \beta^{\min})$. Note that each path p that is not in G_1 can be mapped to an identical path $\hat{p}$ in $\hat{G}$ such that $\mathrm{cost}_{\hat{f}}(\hat{p}) = \mathrm{cost}_f(p)$ and $\mathrm{cost}_{\hat{f}}^+(\hat{p}) = \mathrm{cost}_f^+(p)$. Moreover, each path p that is in G_1 can be mapped to a path $\hat{p}$ in $\hat{G}_1$ coinciding with the (u, t)-subpath (resp. the (s, u)-subpath) p' of p. It can be checked that $\mathrm{cost}_{\hat{f}}(\hat{p})$ and $\mathrm{cost}_{\hat{f}}^+(\hat{p})$ in $\hat{G}$ are obtained by multiplying $\mathrm{cost}_f(p')$ and $\mathrm{cost}_f^+(p')$ in G by the constant $\frac{c(G)}{c(G_1^{ut})}$ (resp. $\frac{c(G)}{c(G_1^{su})}$).

Our goal is to prove that in $\{\hat{p}^1, \ldots, \hat{p}^N\}$ no player has an incentive to deviate. First, we consider the case where $\hat{p}^i$ is in $\hat{G}_1$. Consider the corresponding path p^i chosen by player i in G_1. Since $\{p^1, \ldots, p^N\}$ is a PNE in G, player i cannot improve her cost by deviating to another (u, t)-path in G_1^{ut} (resp. to another (s, u)-path in G_1^{su}). Consequently, player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s, t)-path in $\hat{G}_1$. This implies that in $\hat{G}$

$$\mathrm{cost}_{\hat{f}}(\hat{p}^i) \leq c(\hat{G}_1).$$

Now we show that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s, t)-path outside $\hat{G}_1$. Note that, if we applied Operation 1, we have

$$c(\hat{G}_1) = \beta^{\min}(1 + \alpha^{\min})c(G_1^{ut}) = \frac{c(G)}{c(G_1)}\frac{c(G_1)}{c(G_1^{ut})}c(G_1^{ut}) = c(G).$$

Similarly, if we applied Operation 2, we have

$$c(\hat{G}_1) = \beta^{\min}\left(1 + \frac{1}{\alpha^{\min}}\right)c(G_1^{su}) = \frac{c(G)}{c(G_1)}\frac{c(G_1)}{c(G_1^{su})}c(G_1^{su}) = c(G).$$

21

Clearly, we have $c(\hat{G}_1) = c(G) \leq c(G \setminus G_1) = c(\hat{G} \setminus \hat{G}_1)$. Thus, we obtain

$$\text{cost}_{\hat{f}}(\hat{p}^i) \leq c(\hat{G}_1) \leq c(\hat{G} \setminus \hat{G}_1).$$

We remark that for each path $p \in \hat{G} \setminus \hat{G}_1$ the cost that player i would incur by deviating to p is $\text{cost}_{\hat{f}}^{+}(p)$. Since $\text{cost}_{\hat{f}}^{+}(p) \geq c(\hat{G} \setminus \hat{G}_1)$, we conclude that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path outside $\hat{G}_1$.

Now we consider the case where $\hat{p}^i$ is not in $\hat{G}_1$. Since in G player i cannot improve her cost by deviating from p^i to another (s,t)-path outside G_1 and $G \setminus G_1 = \hat{G} \setminus \hat{G}_1$, we have that in $\hat{G}$ player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path outside $\hat{G}_1$. This also implies that in G $\text{cost}_f(p^i) \leq c(G \setminus G_1)$.

Now we show that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path inside $\hat{G}_1$. Since in G player i cannot improve her cost by deviating from p^i to another (s,t)-path inside G_1, we also have that in G $\text{cost}_f(p^i) \leq c(G_1)$. Thus

$$\text{cost}_f(p^i) \leq c(G) = c(\hat{G}_1).$$

First, recall that $\text{cost}_{\hat{f}}(\hat{p}^i)$ in $\hat{G}$ is equal to $\text{cost}_f(p^i)$ in G. Secondly, note that for each path $p \in \hat{G}_1$ the cost that player i would incur by deviating to p is $\text{cost}_{\hat{f}}^{+}(p)$. Since $c(\hat{G}_1) \leq \text{cost}_{\hat{f}}^{+}(p)$, we conclude that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path inside $\hat{G}_1$. $\diamond$

Claim 4. *The decomposition of $\hat{f}$ into $\{\hat{p}^1, \ldots, \hat{p}^N\}$ obtained by applying Operation 1 (resp. Operation 2) with $(\alpha, \beta) = (\alpha^{max}, \beta^{max})$ is a PNE in the network congestion game on $\hat{G}$.*

Proof of claim. Suppose we apply Operation 1 (resp. Operation 2) with $(\alpha, \beta) = (\alpha^{\max}, \beta^{\max})$. The proof is similar to the previous case, and we will only highlight the main differences. In this case, applying either Operation 1 or Operation 2 yields

$$C(\hat{P}_1) = C(P^*). \tag{25}$$

If $\hat{p}^i$ is in $\hat{G}_1$, player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path in $\hat{G}_1$. Moreover, we have

$$\text{cost}_{\hat{f}}(\hat{p}^i) \leq C(\hat{P}_1) \tag{26}$$
$$= C(P^*) \tag{27}$$
$$\leq c(G) \tag{28}$$
$$\leq c(G \setminus G_1) \tag{29}$$
$$= c(\hat{G} \setminus \hat{G}_1), \tag{30}$$

where (26) follows from the definition of $C(\hat{P}_1)$, (27) follows from (25), (28) follows from the fact that $\{p^1, \ldots, p^N\}$ is a PNE in G, and (30) follows from the definition of $c(G)$.

Since player i would pay $\mathrm{cost}_{\hat{f}}^{+}(p)$ to deviate to a path $p \in \hat{G} \setminus \hat{G}_1$, and because $c(\hat{G} \setminus \hat{G}_1) \leq \mathrm{cost}_{\hat{f}}^{+}(p)$, we conclude that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path outside $\hat{G}_1$.

Now we consider the case where $\hat{p}^i$ is not in $\hat{G}_1$. First, player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path outside $\hat{G}_1$. Moreover, we have

$$\mathrm{cost}_{\hat{f}}(\hat{p}^i) \leq C(P^*) \tag{31}$$

$$= C(\hat{P}_1) \tag{32}$$

$$\leq c(\hat{G}_1), \tag{33}$$

where (31) follows from the fact that $\hat{p}^i = p^i$ and the definition of $C(P^*)$, while (32) follows from (25). Finally, if (33) does not hold, there is a path p in $\hat{G}_1$ such that $\mathrm{cost}_{\hat{f}}^{+}(p) < \mathrm{cost}_{\hat{f}}(\hat{p}_h)$, where $\hat{p}_h$ is the most expensive path in $\{\hat{p}^1, \ldots, \hat{p}^N\}$ that is in $\hat{G}_1$. This would directly imply that player h in G could improve her cost by deviating from p_h by selecting the cheapest path between u and t (resp. s and u), contradicting the fact that $\{p^1, \ldots, p^N\}$ is a PNE in G. $\diamond$

Finally, we prove that also (iii) and (iv) are satisfied if we apply Operation 1 (resp. Operation 2) with appropriate parameters.

Claim 5. The set $\hat{\mathcal{C}} = \{(\hat{C}_i^-, \hat{C}_i^+) : i \in [h]\}$ obtained by applying Operation 1 (resp. Operation 2) with $(\alpha, \beta) = (\alpha^{min}, \beta^{min})$ or $(\alpha, \beta) = (\alpha^{max}, \beta^{max})$ is a collection of pairs of internally disjoint paths in $\hat{G}$ with $|\{\hat{C}_i^- : e \in \hat{C}_i^-\}| \leq \hat{f}_e$ for all $e \in E(\hat{G})$ and such that $\gamma(\hat{\mathcal{C}}^-, \hat{G}) \leq \bar{\gamma}$.

Proof of claim. By construction, the set $\hat{\mathcal{C}}$ obtained with Operation 1 (resp. Operation 2) is a collection of pairs of internally disjoint paths in $\hat{G}$. By construction for each $e \in E(\hat{G})$ we have

$$|\{\hat{C}_i^- : e \in \hat{C}_i^-\}| = |\{C_i^- : e \in C_i^-\}| \leq f_e = \hat{f}_e.$$

We now prove that $\gamma(\hat{\mathcal{C}}^-, \hat{G}) \leq \bar{\gamma}$. To this purpose, we need to decide how to "cover" each $\hat{C}^- \in \hat{\mathcal{C}}^-$ in the expression defining $\gamma(\hat{\mathcal{C}}^-, \hat{G})$.

Suppose we performed Operation 1 (resp. Operation 2).

(a) if $\hat{C}^- = C^-$ for some $C^- \in \mathcal{C}^- \setminus \mathcal{C}^-(G_1)$, we use the path that covered C^- in (16).

(b) if $\hat{C}^- = C^-$ for some $C^- \in \mathcal{C}^-(G_1^{ut})$ (resp. $C^- \in \mathcal{C}^-(G_1^{su})$), we use the subpath from u to t (resp. from s to u) of the path that covered C^- in (16).

(c) if $\hat{C}^- = C_{ut}^-$ (resp. $\hat{C}^- = C_{su}^-$), for some (s,t)-path $C^- \in \mathcal{C}(G_1)$ whose subpath from u to t is C_{ut}^- (resp. whose subpath from s to u is C_{su}^-), we conclude that $\hat{C}^-$ is an (s,t)-path in $\hat{G}$ and we use a copy of $\hat{C}^-$ in (16).

Consider the cycle $C_j^- \in \mathcal{C}^-$ that we used to decompose G^1 and suppose that $C_j^- \in G_1^{ut}$. If we performed Operation 1, (b) implies $\gamma(\hat{\mathcal{C}}^-, \hat{G}) < \gamma(\mathcal{C}^-, G) = \bar{\gamma} + 1$. If we performed

23

Operation 2, C_j^- does not belong to $\hat{\mathcal{C}}$, since it was shrunk during Operation 2. Also in this case $\gamma(\hat{\mathcal{C}}^-, \hat{G}) < \gamma(\mathcal{C}^-, G) = \bar{\gamma} + 1$. If $C_j^- \in G_1^{su}$, we reach the same conclusion by applying similar arguments. $\diamond$

By Claims 3 and 4, and since $\gamma(\hat{\mathcal{C}}^-, \hat{G}) \leq \bar{\gamma}$, we can apply our inductive hypothesis to conclude that $\Delta(\hat{\mathcal{C}}, \hat{f}) \leq \frac{1}{4}\text{cost}(\hat{f})$. Finally, claims 5 and 2 immediately imply $\Delta(\mathcal{C}, f) \leq \frac{1}{4}\text{cost}(f)$. $\square$

5 Proof of Theorem 2

In this section, we provide a lower bound on the PoA of series-parallel network congestion games with affine delays that approaches $27/19$ as $N \to \infty$.

Let $\{q_1, \ldots, q_N\}$ be an ordered sequence of positive numbers such that $\sum_{i=1}^N q_i = 1$ and $q_{i+1} = \frac{1}{2}\sum_{j=1}^i \frac{q_j}{i}$ for $i \in [N-1]$. Let $m \in [N-1]$ and define $r = \frac{m}{N}$. By statement 1 in Lemma 8 and Lemma 9 we have

$$\sum_1^m q_i - \frac{m}{N} = \mu(m, N) - \frac{m}{N} \geq \left[(\sqrt{r} - r) - \epsilon\right],$$

where $\epsilon = \sqrt{r} - \sqrt{\frac{2rN-1}{2N-1}}$.

We define a new sequence $\{s_1, \ldots, s_N\}$ by averaging $\{q_1, \ldots, q_m\}$. Precisely, $s_1 = \cdots = s_m = \frac{\sum_1^m q_i}{m}$ and $s_j = q_j$ for $j \geq m+1$. This implies that

$$\sum_1^m s_i - \frac{m}{N} = \sum_1^m q_i - \frac{m}{N} \geq \left[(\sqrt{r} - r) - \epsilon\right]. \tag{34}$$

Note that $s_{i+1} \geq \frac{1}{2}\sum_{j=1}^i \frac{s_j}{i}$ for $i \in [N-1]$.

We construct a series-parallel (s,t)-network G with affine delays 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 $s_{i+1}x$ 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 $e_1, \ldots, e_m$ with delay function $\frac{1}{N}x$ to get G. By construction, G is a series-parallel network. Fig. 3 illustrates our construction.

Lemma 12 *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 this by induction on $i \in \{m, \ldots, N\}$. 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 $m \leq i < N$ 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
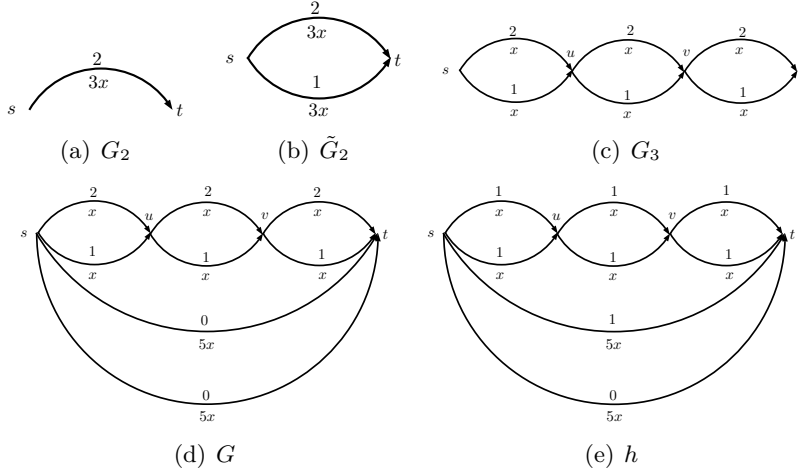
Figure 3: Given an input sequence $\{8, 4, 3\}$ and $m=2$, we first average the first m numbers and get $\{6, 6, 3\}$. (d) is the output network G and its corresponding PNE flow f. (a), (b), (c) are the intermediate networks and flows according to our construction. (e) is the flow h defined in the proof of Theorem 2 where $k = 1$.

(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 $\text{cost}_{f_{i+1}}(\bar{p}^{i+1}) = s_1$. $\qquad \square$

Lemma 13 *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 $2s_{k+1} \geq \frac{1}{k}\sum_{i=1}^k s_i$, 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, $\mathrm{cost}(f_{k+1}) = \mathrm{cost}(\tilde{f}_k) = \sum_{i=1}^{k+1} s_i$. 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. 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 $\frac{1}{N}x$. 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.

$\square$

Proof of Theorem 2. Consider the network congestion game on the network G defined above. By Lemma 12, f has an (s, t)-path $\bar{p}$ with flow value m and $\mathrm{cost}_f(\bar{p}) = s_1$. For each edge e in $\bar{p}$, let $a_e x$ be the delay function of e. Note that $\mathrm{cost}_f(\bar{p}) = \sum_{e \in \bar{p}} a_e m = s_1$ implies that $\sum_{e \in \bar{p}} a_e = \frac{s_1}{m}$. Let $k \in [m]$ and define $l = \frac{k}{m}$. 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 $\frac{1}{N}x$. Then by construction we have:

$$\mathrm{cost}(f) - \mathrm{cost}(h) = m\mathrm{cost}_f(\bar{p}) - \left(k\mathrm{cost}_h(\bar{p}) + (m - k)\frac{1}{N} \right)$$

$$= s_1 m - \left(\frac{s_1}{m}k^2 + (m - k)\frac{1}{N} \right) \tag{35}$$

$$= \left(\frac{s_1}{m}m^2 - \frac{s_1}{m}k^2 - \frac{m - k}{m}ms_1 \right) + \frac{m - k}{m}\left(ms_1 - \frac{m}{N} \right)$$

$$= \left(\frac{s_1}{m}mk - \frac{s_1}{m}k^2 \right) + \frac{m - k}{m}\left(\sum_{i=1}^{m} s_i - \frac{m}{N} \right), \tag{36}$$

where equality (35) holds since $\sum_{e \in \bar{p}} a_e = \frac{s_1}{m}$. Equality (36) holds since the first m of s_i are equal. Now observe that

$$\frac{s_1}{m}m^2 = ms_1 = \frac{m}{N} + \left(\sum_1^m s_i - \frac{m}{N} \right) \geq r + \left[(\sqrt{r} - r) - \epsilon \right] = (\sqrt{r} - \epsilon), \tag{37}$$

where the inequality follows from (34). This implies

$$\frac{s_1}{m}mk - \frac{s_1}{m}k^2 = (l - l^2)\frac{s_1}{m}m^2 \geq (l - l^2)(\sqrt{r} - \epsilon), \tag{38}$$

where the inequality follows from (37).

From (36) and (38) we obtain

$$\mathrm{cost}(f) - \mathrm{cost}(h) \geq (l - l^2)(\sqrt{r} - \epsilon) + (1 - l)\left(\sum_1^m s_i - \frac{m}{N} \right)$$

$$\geq (l - l^2)(\sqrt{r} - \epsilon) + (1 - l)\left[(\sqrt{r} - r) - \epsilon \right], \tag{39}$$

where inequality (39) follows from (34). By Lemma 13 we know that $\text{cost}(f) = \sum_1^N s_i = 1$, thus we obtain:

$$\text{cost}(h) \leq 1 - (l - l^2)(\sqrt{r} - \epsilon) - (1 - l)\left[(\sqrt{r} - r) - \epsilon\right]$$
$$= 1 + l^2\sqrt{r} - rl - \sqrt{r} + r + (1 - l^2)\epsilon \tag{40}$$

To obtain an upper bound on $\text{cost}(h)$ we minimize the right-hand-side of (40) with respect to r and l. Observe that $\epsilon \to 0$ when $N \to \infty$, thus we solve

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

which is achieved at $r = \frac{4}{9}$ and $l = \frac{1}{3}$. Since $\frac{\text{cost}(f)}{\text{cost}(o)} \geq \frac{\text{cost}(f)}{\text{cost}(h)}$, we obtain a lower bound for the PoA that asymptotically approaches $\frac{27}{19}$.

$\square$

We point out that the instance with PoA 15/11 provided by Fotakis in [15] can be obtained with our approach with $N = 3$, $m = 2$ and $k = 1$, see Figures 1 and 3. The crucial insight for obtaining our improved lower bound is that m and k are not fixed a priori, but they are used as parameters. By optimizing over these parameters in the final steps of the proof of Theorem 2, we can achieve our improved lower bound on the PoA.

6 Conclusion

We considered series-parallel network congestion games with affine delays. We have exploited the assumptions on the network topology and delay functions to improve the best known bounds on the PoA. Specifically, we have reduced the upper bound of 5/2, valid for general networks [10], to 2, and we have increased the lower bound of 15/11 provided by Fotakis [15] to 27/19. It remains open whether this gap can be closed or further reduced. We conjecture that our upper bound is *not* tight. In fact, to prove that the PoA is at most 2, we used inequality (1) together with Corollary 4. In particular, (1) is derived by using the upper bound

$$\sum_{e:f_e > o_e} f_e d_e(f_e) \leq \text{cost}(f),$$

while Corollary 4 establishes the upper bound:

$$\Delta(f, o) \leq \frac{1}{4}\text{cost}(f).$$

However, we could find no example where both these upper bounds are simultaneously tight.

Finally, to extend our upper bound on the PoA of series-parallel network congestion games with affine delays to the case where the edge delays belong to the family $\mathcal{P}_d$ of polynomials of degree at most d, one would need to extend the result in Corollary 4 and prove

$$\Delta(f, o) \leq \gamma(\mathcal{P}_d)\text{cost}(f),$$

27

where $\gamma(\mathcal{P}_d)$ is a function of d. A straightforward extension of our approach implies $1 - \beta(\mathcal{P}_d) - \gamma(\mathcal{P}_d) \leq 0$, which, according to (1) leads to an inconsequential upper bound. Thus, a different approach is needed when dealing with polynomial delays.

Acknowledgements

References

[1] Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. J. ACM **55**(6), 25:1–25:22 (2008)

[2] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: 45th Annual IEEE Symposium on Foundations of Computer Science, pp. 295–304 (2004)

[3] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM Journal on Computing **38**(4), 1602–1623 (2008)

[4] 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, STOC '05, p. 57–66. Association for Computing Machinery, New York, NY, USA (2005)

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

[6] Bern, M., Lawler, E., Wong, A.: Linear-time computation of optimal subgraphs of decomposable graphs. Journal of Algorithms **8**(2), 216 – 235 (1987)

[7] Bhaskar, U., Fleischer, L., Huang, C.: The price of collusion in series-parallel networks. In: F. Eisenbrand, F.B. Shepherd (eds.) Integer Programming and Combinatorial Optimization, pp. 313–326. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)

[8] Bilò, V., Vinci, C.: The price of anarchy of affine congestion games with similar strategies. Theoretical Computer Science **806**, 641 – 654 (2020)

[9] 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, STOC '05, p. 67–73. Association for Computing Machinery, New York, NY, USA (2005)

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

[11] Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Selfish routing in capacitated networks. Mathematics of Operations Research **29**, 2004 (2003)

[12] Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: On the inefficiency of equilibria in congestion games. In: Integer Programming and Combinatorial Optimization, pp. 167–181. Springer Berlin Heidelberg, Berlin, Heidelberg (2005)

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

[14] Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: Proceedings of STOC '04 (2004)

[15] Fotakis, D.: Congestion games with linearly independent paths: Convergence time and price of anarchy. In: B. Monien, U.P. Schroeder (eds.) Algorithmic Game Theory, pp. 33–45. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)

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

[17] Holzman, R., Monderer, D.: Strong equilibrium in network congestion games: increasing versus decreasing costs. International Journal of Game Theory **44**, 647–666 (2015)

[18] Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How easy is local search? Journal of Computer and System Sciences **37**(1), 79 – 100 (1988)

[19] Kikuno, T., Yoshida, N., Kakuda, Y.: A linear algorithm for the domination number of a series-parallel graph. Discrete Applied Mathematics **5**(3), 299 – 311 (1983)

[20] Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: STACS 99, pp. 404–413. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)

[21] Krumke, S.O., Zeck, C.: Generalized max flow in series–parallel graphs. Discrete Optimization **10**(2), 155–162 (2013)

[22] Milchtaich, I.: Network topology and the efficiency of equilibrium. Games and Economic Behavior **57**(2), 321–346 (2006)

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

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

[25] Radzik, T.: Faster algorithms for the generalized network flow problem. Math. Oper. Res. **23**(1), 69–100 (1998)

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

[27] Rosenthal, R.W.: The network equilibrium problem in integers. Networks **3**(1), 53–59 (1973)

[28] Roughgarden, T.: 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

[29] Roughgarden, T.: Selfish routing with atomic players. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05, p. 1184–1185. Society for Industrial and Applied Mathematics, USA (2005)

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

[31] Takamizawa, K., Nishizeki, T., Saito, N.: Linear-time computability of combinatorial problems on series-parallel graphs. J. ACM **29**(3), 623–641 (1982)

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