INEFFICIENCY OF THE PURE NASH EQUILIBRIA IN STRUCTURED CONGESTION GAMES

by

Bainian Hao

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Industrial Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2024

Date of final oral examination: 08/12/2024

The dissertation is approved by the following members of the Final Oral Committee:

Carla Michini, Assistant Professor, Industrial and Systems Engineering
Alberto Del Pia, Associate Professor, Industrial and Systems Engineering
Jeffrey Linderoth, Professor, Industrial and Systems Engineering
Victor Zavala, Professor, Chemical and Biological Engineering

To my dearest grandmother, Guicun Wang, whose love have been the light in my journey.

ACKNOWLEDGMENTS

First and foremost, I want to express my deepest gratitude to my advisor, Prof. Carla Michini, for her invaluable support, unwavering patience, and meticulous guidance throughout my Ph.D. journey. Carla has been instrumental in my academic development, from refining my proof-writing skills to shaping high-level research ideas. Moreover, she has been an exceptional role model for me. Her remarkable kindness and optimism have continually inspired and encouraged me, even during my most challenging and unproductive moments. Without Carla's guidance, I would not have been able to achieve any of my accomplishments.

I am also deeply grateful to my committee members, Profs. Alberto Del Pia, Jeffrey Linderoth, and Victor Zavala. Alberto, an exceptional scholar, has continuously inspired me with his brilliant ideas and guided me toward new research directions. Working with Jeff as a teaching assistant has left me profoundly impressed by his outstanding character, and I greatly appreciate the supportive conversations we had during my stressful job search. I also want to extend my heartfelt thanks to Victor for his constructive advice during both my preliminary exam and final defense.

I would also like to extend special thanks to Profs. Yuri Faenza and Anush Tserunyan, who introduced me to the world of optimization and mathematics. Their tremendous support and recommendations were invaluable during my Ph.D. application process.

My time as a student at the University of Wisconsin-Madison has been an unforgettable and profoundly enriching experience. I want to thank Profs. Jelena Diakonikolas, Michael Ferris, Kirthevasan Kandasamy, James Luedtke, Robert Nowak, and Christos Tzamos for their exceptional lectures on computer science and optimization. I also want to express my heartfelt thanks to the friends I met in Madison who made this journey truly special—particularly Katherine B. Adams, Qiulin Chen, Rui Chen, Qinyang He, Yepeng Jin, Qiaomei Li, Zhichao Ma, Ari Smith, Qingwei Wang, Yongjia Wang, Yixuan Zhang, Dekun Zhou, Zachary Zhou, and Haoran Zhu. While it's impossible to name everyone, the camaraderie, encouragement, and shared experiences with my fellow friends have been a constant source of strength and joy throughout this challenging yet rewarding period.

Finally, I want to give my warmest thanks to my parents and all of my family. Your love, patience, and constant encouragement provided me with the strength I needed to persevere through the challenges and pressures of this experience. I couldn't have reached this milestone without your sacrifices and endless belief in me, and I am forever thankful for everything you've done to help me succeed.

CONTENTS

LIST OF FIGURES

ABSTRACT

The Price of Anarchy (PoA), which is the largest ratio between the cost of pure Nash equilibria and the cost of the social optimum, is a classic measure of inefficiency in game theory. In this thesis, we study the PoA of congestion games over special combinatorial structures, specifically series-parallel networks and matroids. Our results show that the inefficiency of equilibria can significantly decrease when players' strategy sets are restricted to the above structures.

We first 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 $15/11$.

Next, we study the PoA of congestion games in series-parallel networks with an arbitrary cost function class $\mathcal{D}$. We introduce a quantity $y(\mathcal{D})$ to upper bound the PoA for delay functions in class $\mathcal{D}$. When $\mathcal{D}$ is the class of polynomial functions with highest degree p, our upper bound is $2^{p+1} - 1$, which is significantly smaller than the worst-case PoA for general networks. Thus, restricting to series-parallel networks can limit the inefficiency of pure Nash equilibria. We also construct a family of instances with polynomial delay functions that have a PoA in $\Omega(2^p/p)$ when the number of players goes to infinity. Compared with the subclass of extension-parallel networks, whose worst-case PoA is in $\Theta\left(p/\ln p\right)$, our results show that the worst-case PoA quickly degrades from sub-linear to exponential when relaxing the network topology. We finally consider an alternative measure of the social cost

of a strategy profile as the maximum players' cost. We introduce a parameter $z(\mathcal{D})$ and we show that the PoA is at most $y(\mathcal{D})z(\mathcal{D})$, which for polynomial delays of maximum degree p is at most 2^{2p}. Compared to the PoA in general networks, which is in $p^{\Theta(p)}$, our results shows a significant improvement in efficiency. We also prove that our previous lower bound in $\Omega(2^p/p)$ is still valid for this measure of social cost. This is in stark contrast with the PoA in the subclass of extension-parallel networks, where each pure Nash equilibrium is a social optimum.

Finally, we study the PoA of congestion games in matroids. We derive new upper bounds on the PoA of symmetric congestion games defined over k-uniform matroids and paving matroids with delay functions in arbitrary class $\mathcal{D}$. Specifically, when $\mathcal{D}$ is the class of polynomial functions with highest degree p, we showed that the upper bounds on the PoA of k-uniform matroids and paving matroids congestion games are in $O(2^p p/\ln p)$ and $O(4^p p/\ln p)$, respectively. Moreover, we also show that for symmetric paving matroid congestion games with affine delay functions, the PoA is at most $17/7$.

1 INTRODUCTION

With the rapid development of the Internet, traditional markets have undergone dramatic transformations and accelerated growth. The digital era has brought a paradigm shift, redefining how markets operate on a global scale. When analyzing these expansive and dynamic markets, scientists face the challenge of not only optimizing the strategic decision-making of individual entities but also understanding and predicting the strategic interactions among multiple agents. This complexity arises because market participants are numerous and diverse, ranging from individuals and small businesses to large corporations and governmental bodies, all of whom interact in ways that significantly impact one another and the overall market dynamics.

The foundation for understanding such interactions can be traced back to the pioneering work of John von Neumann and Oskar Morgenstern. In their seminal 1944 book, *Theory of Games and Economic Behavior*, they laid down the initial ideas of what we now know as modern game theory. Game theory provides a robust framework to articulate, analyze, and eventually comprehend the intricate interplay between agents and their decisions. Essentially, a game is characterized by a set of players who interact according to specific rules. These players engage in strategic decision-making processes that are interdependent, meaning the actions of one player can significantly influence the outcomes for others.

In the context of large-scale networks, such as the Internet and road transportation systems, the enormity and complexity of these networks mean that players often cannot cooperate directly. Instead, they act in their own self-interest, aiming to maximize their utility. Thus, the lack of a central regulation becomes a realistic assumption [31, 46]. This assumption aligns with the principles of non-cooperative games. In non-cooperative games, players' actions are influenced by the actions of others, and the overall quality of each player's strategy is contingent upon the strategies adopted by their counterparts. Moreover, players are assumed to be self-interested and rational, seeking to achieve the best possible outcome given their preferences and the strategies of others.

The concept of Nash equilibrium, named after mathematician John Nash, is a cornerstone in predicting the outcomes of non-cooperative games. At Nash equilibrium, no player can unilaterally change their strategy to achieve a better payoff, given the strategies of all other players. This equilibrium represents a state of stability where players' strategies are in mutual best responses.

In the field of algorithmic game theory, one of the most crucial questions revolves around the existence and computation of Nash equilibria. John Nash's classic theorem [45], established in 1950, assures us that any game with a finite number of players and strategies has at least one mixed Nash equilibrium, where each player chooses a probability distribution over their set of strategies. However, the existence of pure Nash equilibria, where each player selects a single strategy, is not always guaranteed.

In 1973, Robert Rosenthal showed that the pure Nash equilibrium always exist in a class of games, which includes congestion games [50]. This result was later generalized by Monderer and Shapley to the *Potential Games* [44], a significant class of games in which pure Nash equilibria are always present. In these games, players' incentives can be captured by a potential function, ensuring that any improvement by an individual player corresponds to a corresponding improvement in the potential function. This characteristic simplifies the analysis and computation of equilibria, providing a robust framework for understanding strategic interactions in a variety of settings.

Once the existence of equilibrium is assured and their computability is understood, it would be nature for us to think: *How good are these equilibria?* Unfortunately, the pursuit of individual self-interest in a non-cooperative context often leads to outcomes that are suboptimal from a societal perspective. The inefficiency of these equilibria can be stark, as they may not align with the minimization of overall social costs.

To quantify this inefficiency, the concept of the *Price of Anarchy* (PoA) was introduced by Koutsoupias and Papadimitriou in 1999 [34]. The Price of Anarchy measures the worst-case ratio between the cost of a Nash equilibrium and the minimum possible social cost that could be achieved if players cooperated to optimize the overall outcome. This metric highlights the potential disparity between individually rational behavior and collective optimality, providing insights into the potential inefficiencies that arise in decentralized, self-interested environments.

In this thesis, we study congestion games, a subclass of potential games, and focus on the inefficiency of *pure* Nash equilibria (PNE). We study two combinatorial structures, series-parallel networks and paving matroids, and show that the PoA decreases significantly if the resource sets of congestion games are restricted to them.

1.1 Congestion games

Congestion games are a special sub-class of non-cooperative games, where each player aims at selecting a subset of resources from a set of available resources, but the cost of each resource 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 [31, 46] and represent a simple, yet powerful paradigm for selfish resource sharing.

Definition 1.1. A (atomic and unweighted) *congestion game* consists of a finite set of resources R and N players. For each player $i \in [N]$, there is a strategy set $X^i \subseteq 2^R$ and every strategy $x^i \in X^i$ is a subset of R. For each resource $e \in R$ there is a non-decreasing delay function $d_e(x) : [N] \to \mathbb{R}_{\geq 0}$.

A congestion game is called *symmetric* if $X^1 = X^2 = \cdots = X^N$, otherwise it is *asymmetric*. A *state* of the game is a strategy profile $s = (s^1, \ldots, s^N)$ where $s^i \in X^i$ is the strategy chosen by player i, for $i \in [N]$. The set of states of the game is denoted by $X = X^1 \times \cdots \times X^N$. Given a state s and a resource e, we denote by s_e the number of players choosing e in s.

There are also several variants of congestion games. The first distinction is between atomic and non-atomic congestion games. In *non-atomic* congestion games, there are infinitely many players and each player controls an infinitesimal amount of flow. In contrast, in *atomic* congestion games each player controls a non-negligible amount of flow. The second distinction is between weighted and unweighted games. In *weighted* congestion games, the effect of each player on congestion is proportional to the player's weight, while in *unweighted* congestion games all players have the same effect on congestion.

To measure the congestion level in a given state, we define the costs as follows:

Definition 1.2. Given a state s and a set of resources $P \subseteq R$, the *cost* of P is defined as

$$\text{cost}_s(P) = \sum_{e \in P} \text{cost}_s(e) = \sum_{e \in P} d_e(s_e),$$

and the *cost* of the state s is defined as

$$\text{cost}(s) = \sum_{i \in [N]} \text{cost}_s(s^i) = \sum_{e \in R} s_e d_e(s_e).$$

We usually measure the cost of a state in two ways, the *total cost* of a state s, denoted by $\text{tot}(s)$, is the sum of all players' costs. Clearly $\text{tot}(s)$ coincides with the cost of the state s:

$$\text{tot}(s) = \sum_{i \in [N]} \text{cost}_s(s^i) = \text{cost}(s).$$

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

$$\max(s) = \max_{i \in [N]} \text{cost}_s(s^i).$$

1.1.1 Nash equilibria, social optima and the Price of Anarchy

One of the most important concepts in game theory is that of *pure Nash equilibrium* (PNE), which represents a stable outcome of the game, since no player $i \in [N]$ can improve their cost if they unilaterally changes strategy by selecting a different strategy.

Definition 1.3. A *pure Nash equilibrium* (PNE) is a state $s = (s^1, \ldots, s^i, \ldots, s^N)$ such that, for each $i \in [N]$ we have

$$\text{cost}_s(s^i) \leq \text{cost}_{\tilde{s}}(\tilde{s}^i) \qquad \forall (s^1, \ldots, \tilde{s}^i, \ldots, s^N) \in X.$$

A PNE is optimal solution for each single player. A centeral decision maker (e.g. government or centeral authority) would instead persue a strategy profile maximizing the social welfare, which is defined as follows

Definition 1.4. A *social optimum* with respect to the total cost is a state o such that

$$\text{cost}(o) \leq \text{cost}(s) \qquad \forall s \in X.$$

A *social optimum* with respect to the maximum cost is a state o such that

$$\max(o) \leq \max(s) \qquad \forall s \in X.$$

Recall that although pure Nash equilibria are best for each single player, the social costs of pure Nash equilibria could be much higher compared with the cost of the social optimum. Thus, we consider pure Nash equilibria as inefficient outcomes. To measure the inefficiency of pure Nash equilibria, we use the definition of (pure) Price of Anarchy [34].

Definition 1.5. The *Price of Anarchy* (PoA) is the maximum ratio between the cost of a PNE and the cost of a social optimum.

$$\text{PoA} := \max_{f \text{ is a PNE}, s \in X} \frac{\text{cost}(f)}{\text{cost}(s)}$$

In other words, to compute the PoA we consider the social optimal state and the "worst" PNE, i.e., a PNE whose cost is as large as possible.

We remark that for general game theory models, pure Nash equilibria are not guaranteed to exist. However, every game must have a *mixed Nash equilibrium*, in which each player selects strategies at random and acts to minimize their *expected payoff*. To define such randomized strategies formally, let us enhance the choices of players so each one can pick a probability distribution over their set of possible strategies; such a choice is called a *mixed strategy*. We assume that players independently select strategies using the probability distribution. The independent random choices of players lead to a probability distribution of strategy vectors s. Nash (1951) proved that under this extension, every game with a finite number of players, each having a finite set of strategies, has a Nash equilibrium [45].

Theorem 1.6 ([45])**.** *Any game with a finite set of players and finite set of strategies has a Nash equilibrium of mixed strategies.*

1.1.2 Potential function

Although pure Nash equilibria may not exist in general games, they are guaranteed to exist in every congestion game. In the classical paper [50] Rosenthal proved the following theorem:

Theorem 1.7. *Every congestion game has at least one pure Nash equilibrium.*

In the proof Rosenthal introduced the following *potential function* $\Phi(x) : X \to R_{\geq 0}$:

$$\Phi(x) = \sum_{e \in R} \sum_{i=1}^{x_e} d_e(i).$$

Given a state x and a player i, let x^{-i} denotes the strategies of the other players in x. Also let $\bar{x} = (\bar{x}^i, x^{-i}) \in X$, then we have that

$$\Phi(x) - \Phi(\bar{x}) = \text{cost}_x(x^i) - \text{cost}_{\bar{x}}(\bar{x}^i).$$

Thus, given a pure Nash equilibrium x, for each player i and strategy $\bar{x}^i \in X^i$, we have

$$\Phi(x^i, x^{-i}) \leq \Phi(\bar{x}^i, x^{-i}).$$

Let the neighbourhood of a state $\bar{x}$ be $\mathcal{N}(\bar{x}) = \{(x^i, \bar{x}^{-i}) : \text{for all } i \in [N], x^i \in X^i\}$. $\bar{x}$ is a local minimum of $\Phi(x)$ if $\Phi(\bar{x}) \leq \Phi(x)$ for all $x \in \mathcal{N}(\bar{x})$. Then every local minimum of $\Phi(x)$ is a PNE of the congestion game.

Congestion games have been widely investigated in the literature [55, 57, 51, 20, 54, 9]. 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 [50, 44]. The inefficiency of pure Nash equilibria in congestion games has also been widely studied. For atomic congestion games with affine delays and $N \geq 3$ players, Awerbuch et al. [5, 6] and Christodoulou and Koutsoupias [13] independently provided an upper bound of $5/2$ on the PoA. If the game is symmetric the bound can be improved to $(5N - 2)/(2N + 1)$ [13]. Correa et. al. 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. For asymmetric atomic congestion games having polynomial delay functions with nonnegative coefficients and highest degree p, Christodoulou and Koutsoupias [13] showed that the PoA is in $p^{\Theta(p)}$ (see also [5, 6]). Aland et al. [3] 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

Table 1.1: Summary of related work and of our main results for the worst-case PoA with respect to the total players' costs. In the column "structure", we write "-" when no assumption on the structure of strategy sets is made.

Table 1.1					
Atomic	Symmetric	Delays	Structure	PoA	
N	N	Affine	-	$= \rho(\mathcal{D}) = 4/3$	[52]
N	N	Poly-p	-	$= \rho(\mathcal{D}) \in \Theta(p/\ln p)$	[52]
Y	N	Affine	-	$= 5/2$	[13]
Y	Y	Affine	-	$= \frac{5N-2}{2N+1}$	[13, 14]
Y	Y	Affine	ext-para	$= 4/3$	[25, 26]
Y	Y	Affine	ser-para	$\in [27/19, 2]$	[Chapter 2]
Y	Y	Affine	graphic matroid	$= \frac{5N-2}{2N+1}$	[23]
Y	Y	Affine	k-uniform matroid	≤ 1.41	[18]
Y	Y	Affine	paving matroid	$\leq 17/7$	[Chapter 4]
Y	N	Poly-p	-	$\in [\lfloor \Phi_p \rfloor^{p+1}, \Phi_p^{p+1}]$	[3, 4]
Y	Y	Poly-p	-	$\geq \lfloor \Phi_p \rfloor^{p+1} \in \Theta(p/\ln p)^{p+1}$	[Chapter 3]
Y	Y	Poly-p	ext-para	$= \rho(\mathcal{D}) \in \Theta(p/\ln p)$	[25, 26]
Y	Y	Poly-p	ser-para	$\leq y(\mathcal{D}) = 2^{p+1} - 1$	[Chapter 3]
Y	Y	Poly-p	ser-para	$\in \Omega(\frac{2^p}{p})$	[Chapter 3]
Y	Y	Poly-p	k-uniform matroid	$\in O(2^{p(p+1)})$	[36]
Y	Y	Poly-p	k-uniform matroid	$\in 2^p \Theta(p/\ln p)$	[Chapter 4]
Y	Y	Poly-p	paving matroid	$\in 4^p \Theta(p/\ln p)$	[Chapter 4]

ϕ_p^{p+1}, where $\phi_p \in \Theta(p/\ln p)$ is the unique nonnegative real solution to $(x+1)^p = x^{p+1}$.

1.2 Congestion games on networks

1.2.1 Preliminaries on networks

In this section we review the definition of networks and two sub-classes, extension-parallel networks and series-parallel networks. Most of the definitions are from [2].

Definition 1.8. A *directed graph* $G = (V, E)$ consists of a set V of nodes and a set E of arcs whose elements are ordered pairs of distinct nodes. A *directed network* is a directed graph whose nodes and/or arcs have associated numerical values (typically, costs, capacities, and/or supplies and demands).

In this thesis we often make no distinction between "graph" and "network", so we use them synonymously. We also write "directed network" as "network" for convenience. If a network has a single pair of source and sink (s, t), then we call it an (s, t)-network.

Definition 1.9. A graph $G' = (V', E')$ is a *subgraph* of $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E$. We say that $G' = (V', E')$ is the subgraph of G *induced* by V' if E' contains each arc of E with both endpoints in V'.

Definition 1.10. A *walk* in a directed graph $G = (V, E)$ is a subgraph of G consisting of a sequence of nodes and arcs $v_1, e_1, v_2, e_2, \ldots, v_{n-1}, e_{n-1}, v_n$ satisfying the property that for all $1 \leq k \leq n-1$, either $e_k = (v_k, v_{k+1}) \in E$ or $e_k = (v_{k+1}, v_k) \in E$. A *directed walk* is an "oriented" version of a walk in the sense that for any two consecutive nodes v_k, v_{k+1} on the walk, $e_k = (v_k, v_{k+1}) \in E$.

Definition 1.11. A *directed path* is a directed walk without any repetition of nodes.

For convenience, we write "directed path" as "path". Next we introduce two sub-classes of networks.

Definition 1.12. 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.

Definition 1.13 ([25])**.** 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. An

(s,t)-network is *extension-parallel* if it consists of either: (i) a single edge (s,t), (ii) a single edge and an extension-parallel network composed in series, or (iii) two extension-parallel networks composed in parallel.

We remark that series-parallel networks are a superclass of parallel link networks and extension-parallel networks.

Definition 1.14. Given a directed graph $G=(V,E)$ and $s,t \in V$, a vector $x \in \mathbb{R}^E_{\geq 0}$ is called an $(s,t)-$flow if it satisfies:

$$\sum_{(u,v)\in E} x_{u,v} - \sum_{(v,w)\in E} x_{v,w} = 0, \quad \forall v \in V \setminus \{s,t\}.$$

1.2.2 Network congestion games

With the above definitions, we define the *network congestion games* as follows:

Definition 1.15. Let $G=(V,E)$ be a network, we consider a *network congestion game* on G with N players. For each player i, there is a pair of source and sink nodes denoted by (s_i,t_i). The strategy set X^i of player i is the set $\mathcal{P}^i$ of (s_i,t_i)-paths in G. When all the players have the same origin and destination (s,t), their strategy sets all coincide with $\mathcal{P}$ and the game is called *symmetric*.

A *state* of the game is a strategy profile $P=(p^1,\dots,p^N)$ where $p^i \in \mathcal{P}^i$ is the (s_i,t_i)-path chosen by player i, for $i \in [N]$. The set of states of the game is denoted by $X = X^1 \times \cdots \times X^N$. When the game is symmetric, each state $P=(p^1,\dots,p^N) \in X$ 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.

The complexity of PNE in network congestion games has been well studied. Fabrikant et al. [22] 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 for linear delays [22, 1]. Then Del Pia et al. [20] later introduced the class of *totally unimodular* (*TU*) congestion games, where the players' strategies are binary vectors inside polyhedra defined by totally unimodular constraint matrices. Network congestion games belong to this class. In the symmetric case, they gave a strong polynomial algorithm to find a PNE.

There is a rich literature concerning the PoA in network congestion games where the social cost is 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 the network 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 [52] 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 $\mathcal{D}$ [55].

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 [8] 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 [13] showed that the PoA is in $p^{\Theta(p)}$ (see also [5, 6]). Aland et al. [3, 4] later obtained exact values for the worst-case PoA. These exact values admit a lower bound of $\lfloor \Phi_p \rfloor_p^{p+1}$ 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 $p^{\Theta(p)}$ [13, 5, 6]. The worst case PoA drops significantly in the presence of special structure. Lücking et al. [38, 39] studied symmetric congestion games on parallel links and proved that the PoA is $4/3$ for linear functions. Later Fotakis [25, 26] 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 [24, 27]. It is known that, for the class of polynomial delays with nonnegative coefficients and highest degree p, $\rho(\mathcal{D}) \in \Theta\left(p/\ln p\right)$. 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 [30], and it has been shown that the worst-case PoA is between $27/19$ and 2 [30], which is strictly worse than the PoA of $4/3$ in extension-parallel networks [25, 26], and strictly better than the PoA of $5/2$ in general networks [14]. One key step to prove the upper bound in [30] consists in using the following inequality introduced in [25, 26]

$$\frac{\text{cost}(f)}{\rho(\mathcal{D})} \leq \text{cost}(o) + \Delta(f, o), \tag{1.1}$$

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 [30] 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 [30] would provide an inconsequential bound.

Although our focus in the thesis is only on *pure* Nash equilibria, previous works have also investigated the PoA of *mixed* Nash equilibria. For both atomic and non-atomic congestion games in arbitrary networks, the upper bounds on the PoA of pure Nash equilibria also hold for mixed Nash equilibria [3, 4, 52, 54]. In the symmetric case, previous results have showed that the network structure can limit the inefficiency of mixed Nash equilibria. Specifically, Gairing et al. [29] showed that the (mixed) PoA for symmetric unweighted network congestion games on parallel links with polynomial delay functions of highest degree p is at most B_{p+1}, where B_i denotes the Bell number of order i. They further established that this bound is asymptotically tight as the number of players and resources goes to infinity. While the (mixed) PoA bound for congestion games on parallel links is lower than that for arbitrary networks, it exceeds the upper bound presented in Theorem 3.3 for pure Nash equilibria. This implies that the our upper bound on the (pure) PoA of symmetric congestion games over series-parallel networks does not hold for mixed Nash equilibria.

The PoA with respect to the maximum players' cost has received less attention. In the non-atomic setting, Roughgarden [53] showed that the worst-case PoA is $n-1$, where n is the number of nodes in the network.

In the atomic setting, Koutsoupias and Papadimitriou [35] first studied weighted congestion games with linear delay functions on m-parallel links. For the case where the delay functions on the parallel links are all identical, they provided a lower bound in $\Omega\left(\frac{\log m}{\log\log m}\right)$ and an upper bound in $O(\sqrt{m\log m})$ on the worst-case PoA for both pure and mixed Nash equilibria. Later Czumaj and Vöcking [17] extended this framework to parallel links with different linear delay functions, and they established an asymptotically tight bound on the worst-case PoA in $\Theta\left(\frac{\log m}{\log\log\log m}\right)$ for mixed Nash equilibria and in $\Theta\left(\frac{\log m}{\log\log m}\right)$ for pure Nash equilibria. The m-parallel links model has also been studied in [16, 40, 28]. Christodoulou and Koutsoupias [13] investigated the PoA of pure Nash equilibria in general unweighted congestion games. In the symmetric case, they showed that the worst-case 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 worst-case 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. [21] characterized efficient network topologies, i.e., network 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.

1.3 Congestion games on matroids

1.3.1 Preliminaries on matroids

In this section we reivew some basic definitions and properties of matriods. The definitions are from [56].

Definition 1.16. A pair $(S, \mathcal{I})$ is called a matroid if S is a finite set and $\mathcal{I}$ is a nonempty collection of subsets of S satisfying:
(i) if $I \in \mathcal{I}$ and $J \subseteq I$, then $J \in \mathcal{I}$,
(ii) if $I, J \in \mathcal{I}$ and $|I| < |J|$, then $I + z \in \mathcal{I}$ for some $z \in J \setminus I$.

Given a matroid $M = (S, \mathcal{I})$, a subset I of S is called *independent* if I belongs to $\mathcal{I}$, and *dependent* otherwise. A subset B of S is called a *base* of S if B is an inclusionwise maximal independent subset of S. That is, $B \in \mathcal{I}$ and there is no $Z \in \mathcal{I}$ with $B \subset Z \subseteq S$. A set is called simply a *base* if it is a base of S. The common size of all bases is called the *rank* of the matroid. A *circuit* of a matroid $M = (S, \mathcal{I})$ is an inclusionwise minimal dependent set.

Next we introduce three important theorems that characterize the matroid structure from the aspects of bases, circuits and rank functions. The first one is known as the bases exchange property.

Theorem 1.17 ([56])**.** *Let S be a set and let $\mathcal{B}$ be a nonempty collection of subsets of S. Then the following are equivalent:*

1. *$\mathcal{B}$ is the collection of bases of a matroid;*

2. *if $B, B' \in \mathcal{B}$ and $x \in B' \setminus B$, then $B' \setminus \{x\} \cup \{y\} \in \mathcal{B}$ for some $y \in B \setminus B'$;*

3. *if $B, B' \in \mathcal{B}$ and $x \in B' \setminus B$, then $B \setminus \{y\} \cup \{x\} \in \mathcal{B}$ for some $y \in B \setminus B'$.*

In the next theorem, we give the conditions characterizing a collection of circuits of a matroid.

Theorem 1.18 ([56])**.** *Let S be a set and let $\mathcal{C}$ be a collection of nonempty subsets of S, such that no two sets in $\mathcal{C}$ are contained in each other. Then the following are equivalent:*

1. *$\mathcal{C}$ is the collection of circuits of a matroid;*

2. *if $C, C' \in \mathcal{C}$ with $C \neq C'$ and $x \in C \cap C'$, then $(C \cup C') \setminus \{x\}$ contains a set in $\mathcal{C}$;*

3. *if $C, C' \in \mathcal{C}$, $x \in C \cap C'$ and $y \in C \setminus C'$, then $(C \cup C') \setminus \{x\}$ contains a set in $\mathcal{C}$ containing y.*

Finally we give conditions characterizing a rank function of a matroid.

Definition 1.19. The *rank function* of a matroid $M = (S, \mathcal{I})$ is the function $r_M : \mathcal{P}(S) \to \mathbb{Z}_+$ given by:

$$r_M(U) = \max\{|Z| : Z \in \mathcal{I}, Z \subseteq U\}$$

for $U \subseteq S$.

a matroid is determined by its rank function, as a set U is independent if and only if $r(U) = |U|$.

Theorem 1.20 ([56])**.** *Let S be a set and let $r : \mathcal{P}(S) \to \mathbb{Z}_+$. Then r is the rank function of a matroid if and only if for all $T, U \subseteq S$:*

1. *$r(T) \leq r(U) \leq |U|$ if $T \subseteq U$,*

2. *$r(T \cap U) + r(T \cup U) \leq r(T) + r(U)$.*

Next we introduce some classic matroids as examples.

Uniform matroids The easist class of matroids is given by the uniform matroids. They are determined by a set S and a number k: the independent sets are the subsets I of S with $|I| \leq k$. This trivially gives a matroid, called a k-uniform matroid and denoted by U_n^k, where $n := |S|$. In a k-uniform matroid, any subset of S having size greater than or equal to $k+1$ is a circuit.

Graphic matroids Let $G = (V, E)$ be a graph and let $\mathcal{I}$ be the collection of all subsets of E that form a forest. Then $M = (E, \mathcal{I})$ is a matroid. Any spanning forest of G is a base of M, and any circuit in G is a circuit of M as well.

Paving matroids Let $M = (S, \mathcal{I})$ be a matroid has rank t, if all the circuits in M have size at least t, then M is a *paving matroid*.

1.3.2 Matroid congestion games

Given the defition of matroids, then we define the *matroid congestion games* as follows:

Definition 1.21. A *matroid congestion game* consists of N players and a resource set R, such that for each player i the strategy set X^i is the set of bases of a matroid $M_i = (R, \mathcal{I}_i)$. The game is called symmetric if $M_1 = \cdots = M_N$.

For the class of matroid congestion games, a PNE equilibrium can be efficiently computed, both in the symmetric and in the asymmetric case [1, 20]. Concerning the inefficiency of equilibria, Kleer and Schäfer [33] showed that the PoS in general matroids is upper bounded by $\rho(\mathcal{D})$ when the delay functions belong to class $\mathcal{D}$. However, the PoA of matroid congestion games is not well understood. For affine delays, the worst-case PoA of general congestion games, that is equal to $5/2$, can be asymptotically achieved in asymmetric instances of singleton congestion games —that coincide with 1-uniform matroid congestion games— when the number of players goes to infinity [12].

In the symmetric case, the PoA of general matroid congestion games is still not completely understood. For graphic matroids and $N = 2, 3, 4$ or infinity the PoA can be as large as the worst-case PoA of symmetric congestion games, which is equal to $\frac{5N-2}{2N+1}$ [23]. However, for arbitrary N or different delay functions we don't know whether the the worst-case PoA of symmetric congestion games can be achieved by symmetric matroid congestion games. Interestingly, the worst-case PoA of k-uniform matroid congestion games with affine delays cannot exceed 1.4131 and it is equal 1.35188 when the number of players goes to infinity [18]. Moreover, for symmetric k-uniform matroid congestion games with polynomial delays of highest degree p the worst-case PoA is in $O(2^{p(p+1)})$ and in $\Omega(2^p)$ [36]. This indicates that the combinatorial structure of k-uniform matroids significantly limits the inefficiency of equilibria.

1.4 Structure of the thesis

In Chapter 2, we study the Price of Anarchy of congestion games on the series-parallel networks with linear cost functions. We find that the PoA of series-parallel network congestion games is between 27/19 and 2. The upper bound of this PoA is smaller than the PoA of general network congestion games, which is 5/2 [14]. And it is higher than the PoA of extension-parallel network congestion games, which is a subclass of series-parallel congestion games and having PoA at most 4/3 [25]. Motivated by Fotakis [25], for a PNE (s,t)-flow f and a (s,t)-flow o that minimizing the total cost, we look at the function $\Delta(f,o)$ which measures "how different" these two 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 $1/4$ the cost of f, which will imply that the PoA is at most 2. Then we also construct an instance with linear delay functions that have a Price of Anarchy 27/19 as the number of players goes to infinity.

In Chapter 3, we study the Price of Anarchy of congestion games on the series-parallel networks with an arbitrary cost function class $\mathcal{D}$. When $\mathcal{D}$ is the class of polynomial functions with highest degree p, our upper bound is $2^{p+1} - 1$, which is significantly smaller than the worst-case Price of Anarchy in $\Theta(p/\ln p)^p$ for general networks. The main technical ingredient consists in introducing the new parameter $y(\mathcal{D})$ defined in equation (3.1), which intuitively can be used to upper bound by what percentage the cost of an edge increases when one more player uses the edge. First, we assume that the difference $o - f$ contains one cycle including the source s and the sink t. In the crucial step of our proof (Lemma 3.9), we contemplate adding one unit of flow on an arbitrary (s,t)-path p contained in o, and we establish that the corresponding increase of the total cost is at least equal to the average players' cost in f. Next, in Lemma 3.10, we use the definition of $y(\mathcal{D})$ to prove that each (s,t)-path contained in o has cost at least $y(\mathcal{D})$ times the average players' cost in f. Then, if we consider all the (s,t)-paths of an arbitrary decomposition of o, we obtain that cost$(f) \leq y(\mathcal{D})$cost(o) (Lemma 3.11). We remark that this result holds regardless of the network structure. Finally, we relax the assumption on the cycles defined by $o - f$, and we prove a laminarity property of series-parallel networks to prove Theorem 3.2, see Lemmas 3.12 and 3.13.

We also construct a family of instances with polynomial delay functions that have a Price of Anarchy in $\Omega(2^p/p)$ when the number of players goes to infinity. Compared with the subclass of extension-parallel networks, whose worst-case Price of Anarchy is in $\Theta\,(p/\ln p)$, our results show that the worst-case Price of Anarchy quickly degrades from sub-linear to exponential when relaxing the network topology.

We finally consider an alternative measure of the social cost of a strategy profile as the maximum players' cost. We introduce a parameter $z(\mathcal{D})$ and we show that the Price of Anarchy is at most $y(\mathcal{D})z(\mathcal{D})$, which for polynomial delays of maximum degree p is at most 2^{2p}. Compared to the PoA in general networks, which is in $p^{\Theta(p)}$, our results shows a significant improvement in efficiency. We also prove that our previous lower bound in $\Omega(2^p/p)$ is still valid for this measure of social cost. This is in stark contrast with the PoA in the subclass of extension-parallel networks, where each pure Nash equilibrium is a social optimum. The approach that we use for bounding the PoA with respect to the maximum players' cost are based on relating the cost of a social optimum with respect to the maximum players' cost to the cost of a social optimum with respect to the total players' cost. For the upper bound stated in Theorem 3.4, the key step consists in introducing the parameter $z(\mathcal{D})$ defined in equation (3.3) and in establishing that, given a PNE, the most expensive path in the PNE has cost no greater than $z(\mathcal{D})$ times the average players' cost in the PNE. To prove the lower bound stated in Theorem 3.5, given an instance Π with PoA α with respect to the total cost, we construct an instance Π' with PoA α with respect to the maximum players' cost. This new instance is obtained by composing in series N copies of the original network and by permuting the players' strategies so that all the players face the same cost in Π'.

At the end of Chapter 3 we also provide a construction of symmetric congestion games on general networks having polynomial delays function with highest degree p that has Price of Anarchy at least $\lfloor \Phi_p \rfloor^p$, which is very close to the tight PoA of asymmetric network congestion games. This result implies that the low PoA of series-parallel network congestion games is caused not by symmetry but series-parallel structure.

In Chapter 4, we study the PoA of symmetric congestion games on paving matroids and k-uniform matroids. For symmetric k-uniform congestion games with delay functions in an arbitrary class $\mathcal{D}$, we prove that the PoA is upper bounded by $z(\mathcal{D})\rho(\mathcal{D})$. Then we generalized this method from k-uniform matroids to its superclass, paving matroids, and show that for symmetric paving matroid congestion games with delay functions in an arbitrary class $\mathcal{D}$, the PoA is at most $z(\mathcal{D})^2\rho(\mathcal{D})$. We also construct an instance of symmetric paving matroid congestion game that achieves a PoA of 13/9 with affine delay functions, which is worse than the upper bound of the PoA in k-uniform congestion games. Finally, we show that for symmetric paving matroid congestion games with affine delay functions, the PoA is at most 17/7.

Our approach is based on representing the "difference" between a PNE f and a social optimum o of a matroid congestion game as a flow on a complete directed graph, whose nodes correspond to the resources. Each unit of flow on arc (r, r') corresponds to a player replacing r with r' in their strategy. The *overloaded* resources (those with more players in f than in the o) act as supply nodes and the *underloaded* resources (those with more players in the o than in the f) act as demand nodes. If every path from supply u to demand v is such that the costs of u and v in the PNE are related through a constant α, then we can establish that the PoA is at most $\alpha\rho(\mathcal{D})$ (Theorem 4.7). When the delay functions are in class $\mathcal{D}$, we can determine values of α for the case where the matroid is k-uniform (Lemma 4.8) or paving (Lemma 4.10). These results allow us to establish Theorems 4.2 and 4.4. Note that our definition of flows generalizes the idea of the "augmenting paths" used by de Jong et al. [18], extending it from k-uniform matroids with affine delay functions to general matroids with delay functions in class $\mathcal{D}$.

For a paving matroid congestion game with affine delays we require a different approach in order to prove Theorem 4.3. Given f, o and the associated flow, we construct another congestion game with two states s and q such that $\frac{\text{cost}(f)}{\text{cost}(o)} \leq \frac{\text{cost}(s)}{\text{cost}(q)}$. We show that s and q and their associated flow satisfy some special properties, which are used to establish that cost(s)/cost$(q) \leq 17/7$ (Theorem 4.12).

2 PRICE OF ANARCHY IN SERIES-PARALLEL NETWORK CONGESTION GAMES WITH LINEAR DELAYS

2.1 Introduction

2.1.1 Overview and main results

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 [31, 46]. 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.

In this chapter, 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. 2.1 for an example.

First, two-terminal series-parallel networks can be recognized in linear-time [59] 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 [7, 32, 58]. 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 2.1. *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 [14]. In contrast, for network congestion games with affine delays in extension-parallel graphs, Fotakis [25] 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* [43], 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 $15/11 > 4/3$ [25]. 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 2.2. *The PoA of series-parallel congestion games with affine delays is at least $27/19 - \epsilon$, where $\epsilon \to 0$ as $N \to \infty$.*

2.1.2 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 $E(f) = \{(u, v) : (e = (u, v) \in E(G) \text{ and } f_e > 0) \text{ or } (e = (v, u) \in E(G) \text{ and } f_e < 0)\}$. Correspondingly, we denote by $G(f)$ the network $(V(G), E(f))$. Note 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 $\{1, \ldots, n\}$.

Example 1. Consider the 3-player series-parallel congestion game with affine delays depicted in Fig. 2.1. The underlying network G and affine delay functions are showed in Fig. 2.1(a). A PNE flow f is represented in Fig. 2.1(b). The players' strategies are the (s, t)-paths $p^1 = (e_1, e_2, e_6)$, $p^2 = (e_1, e_5, e_3)$, and $p^3 = (e_4, e_2, e_3)$. Moreover, the flow o that

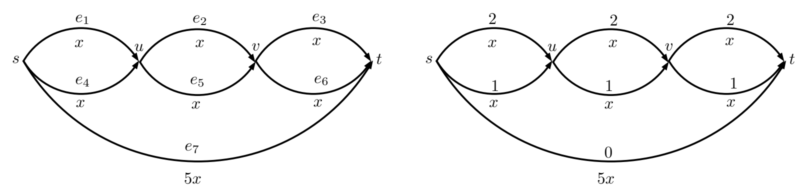

(a) A series-parallel network G (b) A PNE flow f on G

Figure 2.1: The series-parallel network congestion game of Example 1. The PoA is $15/11$.

minimizes the social cost can be reached from f by deviating one unit of flow from the path (e_1, e_2, e_3) to the path (e_7). As a result, we have cost$(f) = 15$ and 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 $4/3$ [25].

2.2 Upper bound on price of anarchy

2.2.1 Overview of the 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 [25], which is tailored to extension-parallel networks.

The first crucial property exploited in [25] 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 [25] 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 [37, 49]. We prove several properties of greedy decompositions and we crucially exploit these properties 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 [25] 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 $1/4$ the cost of f, which will imply that the PoA is at most 2. A similar approach was proposed by [19], who study the PoA in k-uniform matroid congestion games. In particular, the authors show that an analogue of $\Delta(f,o)$, which measures the "difference" between a PNE f and a strategy profile o that minimizes the social cost in the k-uniform matroid congestion game, cannot exceed a constant fraction of the cost of f. We point out that k-uniform matroids and flows in series-parallel networks are very different combinatorial objects, thus the techniques used in [19] cannot be extended to bound $\Delta(f,o)$ in our setting.

2.2.2 Proof of theorem 2.1

To prove Theorem 2.1, we will use the function $\beta(\mathcal{D}) := \sup_{d\in\mathcal{D}} \beta(d)$ introduced in [15], 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 [25]):

$$\text{cost}(f) \leq \text{cost}(o) + \beta(\mathcal{D})\text{cost}(f) + \Delta(f,o). \tag{2.1}$$

If f is a global minimum of the potential function, then $\Delta(f,o) \leq 0$ [25]. However, series-parallel networks might admit PNE that do *not* minimize the potential function. To prove Theorem 2.1, we will exploit the special structure of series-parallel networks and affine delays, in order to show that $\Delta(f,o) \leq \frac{1}{4}\text{cost}(f)$, see Theorem 2.3 and Corollary 2.4. This immediately implies that cost$(f) \leq 2$cost(o), establishing that the PoA is at most 2 for the case under consideration. The main ideas of the proof are described as follows.

To prove Theorem 2.1 we use the following key result.

Theorem 2.3. *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 $|\left\{C_i^- : e \in C_i^-\right\}| \leq f_e$ for all $e \in E(G)$. Then*

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

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 [25]. 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

2.3. 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 2.4. *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 $\Delta(f, o) \leq \frac{1}{4}$cost$(f)$.*

We remark that, to prove Theorem 2.1, it is sufficient to use inequality (2.1) in conjunction with Corollary 2.4.

2.2.3 Proof of theorem 2.3

In this section, we formally prove Theorem 2.3. 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}.

Example 1 (continued)**.** Consider again the network congestion game in Fig. 2.1 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 $\bar{p}^1 = \bar{p}^2 = (e_1, e_2, e_3)$ and $\bar{p}^3 = (e_4, e_5, e_6)$. 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 2.5. *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 $G(f)$. □

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 2.6. *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 $\bar{P} = \left\{\bar{p}^1, \ldots, \bar{p}^N\right\}$ 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

$$R(P, c, x) \leq R(\bar{P}, c, x). \tag{2.2}$$

trivially holds.

Assume that (2.2) holds for $N \leq k$. When $N = k+1$ we first prove that (2.2) 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 $\hat{P} = \left\{\hat{p}^1, \hat{p}^2, \ldots, \hat{p}^N\right\}$ 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

$$\min\{\ell(P') : R(P, c, \bar{x}) \leq R(P', c, \bar{x}),\ P' \text{ decomposition of } f\} \tag{2.3}$$

is zero. Let $\tilde{P} = \{\tilde{p}^1, \ldots, \tilde{p}^N\}$ be a decomposition of f that achieves the minimum in (2.3) 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 $\breve{P}_{uv}$ of h such that $R(\breve{P}_{uv}, c, \bar{x}) \geq R(\tilde{P}_{uv}, c, \bar{x})$ and $\ell(\breve{P}_{uv}) < \ell(\tilde{P}_{uv})$. 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 $\breve{P}_{uv}$ 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 $\breve{P}_{uv}$. 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 $\breve{P}_{uv}$ from $Q \cup \pi$ by replacing π with $\breve{\pi} = \pi \setminus \pi_{uv} \cup q^1_{uv}$ and by replacing q^1 with $q^1 \setminus q^1_{uv} \cup \pi_{uv}$. This immediately implies

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

We prove that

$$R(\breve{P}_{uv}, c, \bar{x}) \geq R(Q \cup \pi, c, \bar{x}). \tag{2.4}$$

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

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

If $c(q^1) \leq \bar{x}$, since $\delta \geq 0$, we immediately have (2.4). If $c(q^1) \geq \bar{x}$ and $c(q^1) - \delta \geq \bar{x}$, we get

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

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 (2.4) 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 (2.2) 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 (2.2) 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

$$R(P', c, x) = \sum_{p \in P'} \max\{0, c(p) - x\} \geq \max\{0, c(f) - (k+1)x\}. \tag{2.5}$$

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 (2.5) $R(\bar{P}, c, x) \geq 0$, thus (2.2) 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 (2.5) $R(\bar{P}, c, x) \geq c(f) - (k+1)x$, thus (2.2) 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

$$c_{e_i} = \begin{cases} \max\{0, \alpha\} & i = 1 \\ \min\{0, \alpha\} & i = k+1 \\ 0 & i = 2, \ldots, k, \end{cases} \tag{2.6}$$

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 (2.2) 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 2.5, $x \leq c(p^1) \leq c(\bar{p}^1)$ we have

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

thus (2.2) 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

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

Since, by Lemma 2.5, $x \geq c(p^{k+1}) \geq c(\bar{p}^{k+1})$ we have

$$R(Q, c', x) = R(P, c, x)$$
$$R(\bar{Q}, c', x) = R(\bar{P}, c, x),$$

thus (2.2) holds.

□

Lemma 2.7. *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 2.7. We need to prove that when $|E| = k+1$, Lemma 2.7 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 $P^* = \{p^1, p^2, \ldots, p^N\}$ 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 $P_1^* = \{p_1^1, p_1^2, \ldots, p_1^N\}$ is a decomposition of f_1, $P_2^* = \{p_2^1, p_2^2, \ldots, p_2^N\}$ 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 $\bar{P}_1 = \{\bar{p}_1^1, \bar{p}_1^2, \ldots, \bar{p}_1^N\}$ and $\bar{P}_2 = \{\bar{p}_2^1, \bar{p}_2^2, \ldots, \bar{p}_2^N\}$ 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 2.7 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

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

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

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

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

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

Note that $\bar{P}^i = \bar{P}^i_1 \cup \bar{P}^i_2$. Our goal is to prove

$$c(\bar{p}^{i+1}) \geq \frac{1}{2}\sum_{j=1}^{i}\frac{c(\bar{p}^j)}{i} = \frac{1}{2}\frac{\sum_{p\in\bar{P}^i_1}c(p) + \sum_{p\in\bar{P}^i_2}c(p)}{|\bar{P}^i_1| + |\bar{P}^i_2|}. \tag{2.7}$$

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

$$c(\bar{p}^{i+1}) \geq \frac{1}{2}\frac{\sum_{p\in\bar{P}^i_1}c(p)}{|\bar{P}^i_1|}. \tag{2.8}$$

If $\bar{P}^i_2 = \emptyset$, (2.7) trivially holds. Thus we now assume $\bar{P}^i_2 \neq \emptyset$. 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

$$c(\bar{p}^{i+1}) \geq \frac{1}{2}\frac{\sum_{p\in\bar{P}^i_2}c(p)}{|\bar{P}^i_2|}. \tag{2.9}$$

Note that (2.8) and (2.9) imply (2.7).

If $|\bar{P}^i_2| = |\bar{P}_2|$, 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:

$$\begin{aligned}
c(\bar{p}^{i+1}) &= \sum_{e\in\bar{p}^{i+1}} d_e(f_e) \\
&\geq \frac{1}{2}\sum_{e\in\bar{p}^{i+1}} d_e(f_e+1) && (2.10)\\
&\geq \frac{1}{2}\max\{c(p) : p \in P^*_2\} && (2.11)\\
&\geq \frac{1}{2}\frac{\sum_{p\in\bar{P}_2}c(p)}{|\bar{P}_2|}. && (2.12)
\end{aligned}$$

Inequality (2.10) 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 (2.11) holds since P^* is a PNE and because $\bar{p}^{i+1}$ is an (s,t)-path. Finally, inequality (2.12) 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 $\ell = \min\left\{t : \bar{p}^t \in \bar{P}_2, t > i+1\right\}$. 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 (2.9). This completes the proof. □

Example 1 (continued)**.** We illustrate the properties stated in Lemmas 2.6 and 2.7 on the example in Fig. 2.1. 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 2.6 holds since $R(P, c, \bar{x}) = 0 < 2 = R(\bar{P}, c, \bar{x})$. For Lemma 2.7, 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 2.7, from $\bar{P}(f) = \left\{\bar{p}^1, \bar{p}^2, \dots, \bar{p}^N\right\}$, we get a sequence of positive numbers $\left\{\text{cost}_f(\bar{p}^1), \text{cost}_f(\bar{p}^2), \dots, \text{cost}_f(\bar{p}^N)\right\}$ 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 2.8. *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 backward 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:

$$\sum_{j=1}^{N-1} x_j = 1 - x_N \leq 1 - \frac{1}{2(N-1)} \sum_{j=1}^{N-1} x_j. \tag{2.13}$$

By equation (2.13), 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 1 holds for $m \in \{k, \dots, N-1\}$, 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:

$$\sum_{j=1}^{k-1} x_j = \sum_{j=1}^{k} x_j - x_k \leq \mu(k, N) - \frac{1}{2(k-1)} \sum_{j=1}^{k-1} x_j. \tag{2.14}$$

According to (2.14), 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 (2.13) and (2.14) with equalities. □

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

Lemma 2.9. *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.

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

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

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

□

From the previous results we can establish the following property of a PNE f, which will be used to prove Theorem 2.3.

Lemma 2.10. *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

$$\sum_{i=1}^{m} c(\bar{p}^i) \leq \left(\frac{1}{4} + \frac{m}{N}\right)\text{cost}(f). \tag{2.15}$$

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 2.7 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 2.8 and 2.9 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 (2.15), 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 $x \in [0,1]$.

$\square$

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

Lemma 2.11. *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 $\text{cost}^+_f(p) \geq \frac{\text{cost}(f)}{N}$.*

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. $\square$

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

Proof of Theorem 2.3. 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

$$\gamma(\mathcal{C}^-, G) = \min \left\{ \sum_{i=1}^{k} |p_i \setminus C_i^-| : p_i \text{ is an } (s,t)\text{-path containing } C_i^-, i \in [k] \right\}. \tag{2.16}$$

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:

$$\begin{aligned}
\Delta(\mathcal{C}, f) &= \sum_{i=1}^{k} (\text{cost}_f(C_i^-) - \text{cost}_f^+(C_i^+)) \\
&\leq \sum_{i=1}^{k} \left(\text{cost}_f(C_i^-) - \frac{\text{cost}(f)}{N} \right) && (2.17) \\
&\leq R\left(\mathcal{C}^-, c, \frac{\text{cost}(f)}{N}\right) && (2.18) \\
&\leq R\left(P, c, \frac{\text{cost}(f)}{N}\right) && (2.19) \\
&\leq R\left(\bar{P}, c, \frac{\text{cost}(f)}{N}\right) && (2.20) \\
&\leq \frac{1}{4}\text{cost}(f). && (2.21)
\end{aligned}$$

Inequality (2.17) 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 2.11, we have $\text{cost}_f^+(C_i^+) \geq \frac{\text{cost}(f)}{N}$ for all $i \in [k]$. Inequality (2.18) follows from the definition of c and the fact that $R\left(\mathcal{C}^-, c, \frac{\text{cost}(f)}{N}\right)$ only contains the nonnegative terms of the summation in (2.17). Inequality (2.19) holds since $\mathcal{C}^- \subseteq P$. Inequality (2.20) holds because of Lemma 2.6. Inequality (2.21) is implied by Lemma 2.10.

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

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

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) $\dfrac{\Delta(\mathcal{C}, f)}{\text{cost}(f)} \leq \dfrac{\Delta(\hat{\mathcal{C}}, \hat{f})}{\text{cost}(\hat{f})}$.

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

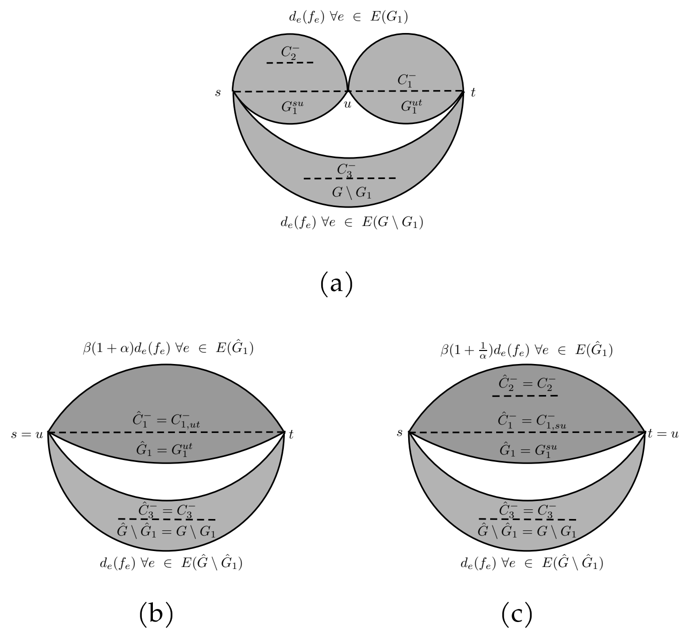

Figure 2.2: This is an example for the operations in Theorem 2.3, where the dotted lines are paths C_i^- in the set $\mathcal{C}^-$. (a) A series-parallel network G. (b) The network $\hat{G}$ by applying Operation 1 to G and (c) the network $\hat{G}$ by applying Operation 2 to G.

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. 2.2.

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 $\text{cost}_f(C_i)$ the difference $\text{cost}_f(C_i^-) - \text{cost}_f^+(C_i^+)$. Let

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

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

Claim 2. If

$$\frac{\Delta(\mathcal{C}, f)}{cost(f)} \leq \frac{D}{cost(f_1)}, \tag{2.22}$$

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

$$\frac{\Delta(\mathcal{C}, f)}{cost(f)} \leq \frac{\Delta(\hat{\mathcal{C}}, \hat{f})}{cost(\hat{f})}. \tag{2.23}$$

Otherwise, if inequality (2.22) *does not hold, for each* $\beta \leq 1$ *and* $\alpha > 0$ *either Operation 1 or Operation 2 yields* (2.23).

Proof of claim. If inequality (2.22) holds and we choose $\beta \geq 1$, then we have:

$$\frac{\Delta(\mathcal{C}, f)}{\text{cost}(f)} \leq \frac{\Delta(\mathcal{C}, f) + (\beta - 1)D}{\text{cost}(f) + (\beta - 1)\text{cost}(f_1)}. \tag{2.24}$$

If inequality (2.22) does not hold and we choose $\beta \leq 1$, then (2.24) still holds. Define

$$\begin{aligned}
B(\xi) &= \xi\text{cost}(f_1^{ut}) - \text{cost}(f_1^{su}) \\
A(\xi) &= -\sum_{C_i^- \in \mathcal{C}^-(G_1^{su})} \text{cost}_f(C_i) + \xi \sum_{C_i^- \in \mathcal{C}^-(G_1^{ut})} \text{cost}_f(C_i) \\
&\quad - \sum_{C_i^- \in \mathcal{C}^-(G_1)(s,t)\text{-path}} \left(\text{cost}_f((C_i^-)_{su}) - \xi\text{cost}_f((C_i^-)_{ut})\right).
\end{aligned}$$

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

$$\begin{aligned}
\Delta(\hat{\mathcal{C}}, \hat{f}) &= \Delta(\mathcal{C}, f) + A(\beta(1+\alpha) - 1) = \Delta(\mathcal{C}, f) + (\beta - 1)D + \beta A(\alpha) \\
\text{cost}(\hat{f}) &= \text{cost}(f) + B(\beta(1+\alpha) - 1) = \text{cost}(f) + (\beta - 1)\text{cost}(f_1) + \beta B(\alpha).
\end{aligned}$$

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

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

Thus, if

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

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

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

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

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

By choosing β appropriately, by (2.24) 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

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

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 $\text{cost}_{\hat{f}}(\hat{p}) = \text{cost}_f(p)$ and $\text{cost}^+_{\hat{f}}(\hat{p}) = \text{cost}^+_f(p)$. 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 $\text{cost}_{\hat{f}}(\hat{p})$ and $\text{cost}^+_{\hat{f}}(\hat{p})$ in $\hat{G}$ are obtained by multiplying $\text{cost}_f(p')$ and $\text{cost}^+_f(p')$ 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 $\hat{G}$

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

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

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

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

$$C(\hat{P}_1) = C(P^*). \tag{2.25}$$

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

$$\begin{aligned}
\text{cost}_{\hat{f}}(\hat{p}^i) &\le C(\hat{P}_1) && (2.26)\\
&= C(P^*) && (2.27)\\
&\le c(G) && (2.28)\\
&\le c(G \setminus G_1) && (2.29)\\
&= c(\hat{G} \setminus \hat{G}_1), && (2.30)
\end{aligned}$$

where (2.26) follows from the definition of $C(\hat{P}_1)$, (2.27) follows from (2.25), (2.28) follows from the fact that $\{p^1, \dots, p^N\}$ is a PNE in G, and (2.30) follows from the definition of $c(G)$. Since player i would pay $\text{cost}^+_{\hat{f}}(p)$ to deviate to a path $p \in \hat{G} \setminus \hat{G}_1$, and because $c(\hat{G} \setminus \hat{G}_1) \le \text{cost}^+_{\hat{f}}(p)$, we conclude that player i cannot improve her cost by deviating from $\hat{p}^i$ to another (s,t)-path 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

$$\begin{aligned}
\text{cost}_{\hat{f}}(\hat{p}^i) &\le C(P^*) && (2.31)\\
&= C(\hat{P}_1) && (2.32)\\
&\le c(\hat{G}_1), && (2.33)
\end{aligned}$$

where (2.31) follows from the fact that $\hat{p}^i = p^i$ and the definition of $C(P^*)$, while (2.32) follows from (2.25). Finally, if (2.33) does not hold, there is a path p in $\hat{G}_1$ such that $\text{cost}^+_{\hat{f}}(p) < \text{cost}_{\hat{f}}(\hat{p}_h)$, 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

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

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 (2.16).

(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 (2.16).

(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 (2.16).

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

2.3 Lower bound on price of anarchy

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 2 in Lemma 2.8 and Lemma 2.9 we have

$$\sum_{i=1}^m q_i - \frac{m}{N} = \mu(m, N) - \frac{m}{N} \geq (\sqrt{r} - r) - \epsilon,$$

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, $s_1 = \cdots = s_m = \frac{\sum_{i=1}^m q_i}{m}$ and $s_j = q_j$ for $j \geq m+1$. This implies that

$$\sum_{i=1}^m s_i - \frac{m}{N} = \sum_{i=1}^m q_i - \frac{m}{N} \geq (\sqrt{r} - r) - \epsilon. \tag{2.34}$$

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. 2.3 illustrates our construction.

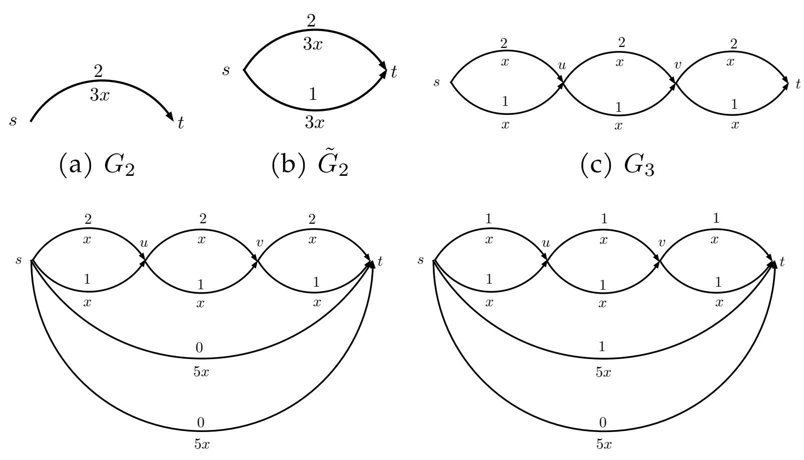

(a) G_2 (b) $\tilde{G}_2$ (c) G_3

(d) G (e) h

Figure 2.3: Consider the input sequence $\{8, 4, 3\}$ and m=2. For convenience we work with integer numbers, but we can easily scale the numbers of the sequence so that they sum up to 1. 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 2.2 where $k = 1$.

Lemma 2.12. *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 (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 $\text{cost}_{f_{i+1}}(\bar{p}^{i+1}) = s_1$. □

Lemma 2.13. *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, $\text{cost}(f_{k+1}) = \text{cost}(\tilde{f}_k) = \sum_{i=1}^{k+1} s_i$. 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. 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 2.2. Consider the network congestion game on the network G defined above. By Lemma 2.12, 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 $a_e x$ 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 $\sum_{e\in\bar{p}} a_e = \frac{s_1}{m}$. 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 $\bar{p}$ 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:

$$
\begin{aligned}
\text{cost}(f) - \text{cost}(h) &= m\text{cost}_f(\bar{p}) - \left(k\text{cost}_h(\bar{p}) + (m-k)\frac{1}{N}\right) \\
&= s_1 m - \left(\frac{s_1}{m}k^2 + (m-k)\frac{1}{N}\right) \qquad (2.35) \\
&= \left(\frac{s_1}{m}m^2 - \frac{s_1}{m}k^2 - \frac{m-k}{m}ms_1\right) + \frac{m-k}{m}\left(ms_1 - \frac{m}{N}\right) \\
&= \left(\frac{s_1}{m}mk - \frac{s_1}{m}k^2\right) + \frac{m-k}{m}\left(\sum_{i=1}^{m} s_i - \frac{m}{N}\right), \qquad (2.36)
\end{aligned}
$$

where equality (2.35) holds since $\sum_{e\in\bar{p}} a_e = \frac{s_1}{m}$. Equality (2.36) holds since the first m of s_i

are equal. Now observe that

$$\frac{s_1}{m}m^2 = ms_1 = \frac{m}{N} + (\sum_{i=1}^{m} s_i - \frac{m}{N}) \geq r + \left[(\sqrt{r} - r) - \epsilon\right] = (\sqrt{r} - \epsilon), \tag{2.37}$$

where the inequality follows from (2.34). This implies

$$\frac{s_1}{m}mk - \frac{s_1}{m}k^2 = (l - l^2)\frac{s_1}{m}m^2 \geq (l - l^2)(\sqrt{r} - \epsilon), \tag{2.38}$$

where the inequality follows from (2.37).

From (2.36) and (2.38)we obtain

$$\begin{aligned}\text{cost}(f) - \text{cost}(h) &\geq (l - l^2)(\sqrt{r} - \epsilon) + (1 - l)(\sum_{i=1}^{m} s_i - \frac{m}{N}) \\ &\geq (l - l^2)(\sqrt{r} - \epsilon) + (1 - l)\left[(\sqrt{r} - r) - \epsilon\right], \end{aligned} \tag{2.39}$$

where inequality (2.39) follows from (2.34). By Lemma 2.13 we know that cost$(f) = \sum_{i=1}^{N} s_i = 1$, thus we obtain:

$$\begin{aligned}\text{cost}(h) &\leq 1 - (l - l^2)(\sqrt{r} - \epsilon) - (1 - l)\left[(\sqrt{r} - r) - \epsilon\right] \\ &= 1 + l^2\sqrt{r} - rl - \sqrt{r} + r + (1 - l^2)\epsilon. \end{aligned} \tag{2.40}$$

To obtain an upper bound on cost(h) we minimize the right-hand-side of (2.40) with respect to r and l. Observe that $\epsilon \to 0$ when $N \to \infty$, thus we solve

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

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 $\frac{27}{19}$.

□

We point out that the instance with PoA $15/11$ provided by Fotakis in [25] can be obtained with our approach with $N = 3$, $m = 2$ and $k = 1$, see Figures 2.1 and 2.3. 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 2.2, we can achieve our improved lower bound on the PoA.

2.4 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 [14], to 2, and we have increased the lower bound of $15/11$ provided by Fotakis [25] 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 (2.1) together with Corollary 2.4. In particular, (2.1) 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 2.4 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 2.4 and prove

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

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 (2.1) leads to an inconsequential upper bound. Thus, a different approach is needed when dealing with polynomial delays.

3 PRICE OF ANARCHY IN SERIES-PARALLEL NETWORK CONGESTION GAMES WITH POLYNOMIAL DELAYS

In this chapter, we still focus on the class of two-terminal series-parallel networks, and we provide upper and lower bounds on the worst-case PoA for (atomic, unweighted, symmetric) network congestion games. Unlike Chapter 2, we consider more general classes of delay functions instead of affine delay functions. Furthermore, except the PoA with respect to the *total cost*, we also study the PoA with respect to the *maximum cost*. Recall that the *total cost* of a state P, denoted by $\mathrm{tot}(P)$, is the sum of all players' costs. Clearly $\mathrm{tot}(P)$ coincides with the cost of the flow $f(P)$:

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

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

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

3.1 Main results

First, we consider the total players' cost. For *asymmetric* network congestion games with polynomial delay functions of highest degree p, Aland et al. provided a tight bound for the worst-case PoA which is in $[\lfloor\Phi_p\rfloor^{p+1}, \Phi_p^{p+1}]$, where $\Phi_p \in \Theta(p/\ln p)$ [3, 4]. Our first goal is to determine whether in *symmetric* network congestion games the same worst-case PoA can be achieved. In Section 3.2, we answer this question in the affirmative.

Theorem 3.1. *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 3.1 indicates that symmetry has little impact on the worst-case PoA of network congestion games.

In Section 3.3.1 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

$$y(\mathcal{D}) = \sup_{d \in \mathcal{D},\ x \in \mathbb{N}^+} \frac{(x+1)d(x+1) - xd(x)}{d(x)}, \tag{3.1}$$

which intuitively can be used to upper bound by what percentage the cost of an arc increases when one more player uses the arc. 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 3.2. *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 3.1) to $O(2^{p+1})$ in symmetric games over series-parallel networks.

In Section 3.3.2 we also provide a lower bound on the worst-case PoA in symmetric congestion games defined over series-parallel networks. This bound is independent of $y(\mathcal{D})$.

Theorem 3.3. *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*

$$\frac{1}{1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r}, \tag{3.2}$$

Table 3.1: Numerical comparison of our bounds on the (pure) PoA of symmetric series-parallel network congestion games with polynomial delays of highest degree p with: (i) the lower bound on the worst-case PoA of symmetric congestion games on arbitrary networks established in Theorem 3.1 (first column); and (ii) the worst-case PoA of symmetric extension-parallel congestion games [25, 26, 27] (last column). The upper bound of 2 for series-parallel networks and affine delays is from [30].

p	LB arbitrary (Theorem 3.1)	UB series-parallel (Theorem 3.2)	LB series-parallel (Theorem 3.3)	PoA ext-parallel [25, 26]
1	1	2 [30]	1.421	1.333
2	8	7	1.938	1.626
3	16	15	2.884	1.896
4	243	31	4.548	2.151
5	729	63	7.491	2.394
6	2187	127	12.747	2.630
7	65536	255	22.228	2.858
8	262144	511	39.48	3.081
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	$\lfloor \Phi_p \rfloor^{p+1} \in \Theta(p/\ln p)^{p+1}$	$2^{p+1}-1$	$\Omega(2^p/p)$	$\rho(\mathcal{D}) \in \Theta(p/\ln p)$

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 $\Theta(p/\ln p)$ [25, 26, 27], our result shows that the PoA dramatically increases when relaxing the network topology from extension-parallel to series-parallel.

In Table 3.1 we provide a numerical comparison of the bounds established in Theorems 3.1, 3.2, 3.3, and the worst-case PoA in extension-parallel networks.

In Section 3.4 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

$$z(\mathcal{D}) = \sup_{d \in \mathcal{D},\ x \in \mathbb{N}^+} \frac{d(x+1)}{d(x)}. \tag{3.3}$$

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

Theorem 3.4. *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 $p^{\Theta(p)}$ [13], 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 3.5. *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 3.5 implies that the worst-case PoA is in $\Omega(2^p/p)$. This is in stark contrast with the result of [21], 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 network operations that we need to allow are the series compositions, which are forbidden for extension-parallel networks.

3.2 Lower bound on symmetric congestion games in general networks

Let $g_p(x) = (x+1)^p - x^{p+1}$. We recall that Φ_p is the unique nonnegative real solution to $g_p(x) = 0$ [3]. Clearly, $g_p(0) = 1 > 0$. Moreover, by Lemma 5.2 in [3], 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 $0 \leq x \leq \phi_p$ we have $g_p(x) \geq 0$ and for $x \geq \phi_p$ we have $g_p(x) \leq 0$. Next, we prove two additional properties of ϕ_p.

Lemma 3.6. *For every positive integer p and for all $k = 0, 1, \ldots, p-1$, we have $(\Phi_{p-1}+1)^k \leq \Phi_{p-1}^{k+1}$.*

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 $k \in \{2, \ldots, p-1\}$, and we prove that the claim also holds for $k-1$. We have

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

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

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

□

Lemma 3.7. *For every positive integer p we have $\Phi_{p-1} \leq \Phi_p \leq \Phi_{p-1} + 1$.*

Proof. We first prove that $\Phi_{p-1} \leq \Phi_p$ 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 $\Phi_p \leq \Phi_{p-1}+1$. 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

$$(\Phi_{p-1}+2)^p < (\Phi_{p-1}+1)^{p+1}. \tag{3.4}$$

We rewrite the left-hand-side of (3.4) as follows.

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

where the first inequality holds by Lemma 3.6. □

Proof of [. Theorem 3.1] 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 $N \geq \max\{\lfloor \Phi_p \rfloor + 1, 2\lfloor \Phi_p \rfloor - 2\}$ 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 $v_{i,0}, v_{i,1}, \ldots, v_{i,\lfloor \Phi_p \rfloor + N-1}$. 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. 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

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 $\lfloor \Phi_p \rfloor \leq j \leq N$ and $c_j x^p$ for $1 \leq j \leq \lfloor \Phi_p \rfloor - 1$ and $N + 1 \leq j \leq \lfloor \Phi_p \rfloor + N - 1$. For each j such that $1 \leq j \leq \lfloor \Phi_p \rfloor - 1$, 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 $N + 1 \leq j \leq \lfloor \Phi_p \rfloor + N - 1$ 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.

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

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

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,

$$\frac{\text{cost}(P)}{\text{cost}(SO)} \geq \frac{\text{cost}(P)}{\text{cost}(P^*)} \tag{3.5}$$
$$= \lim_{N\to+\infty} \frac{(N - \lfloor \Phi_p \rfloor + 1) \cdot \lfloor \Phi_p \rfloor \cdot \lfloor \Phi_p \rfloor^p + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor - 1} c_j \cdot j^p}{(N - \lfloor \Phi_p \rfloor + 1) + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor - 1} c_j} \tag{3.6}$$
$$= \lim_{N\to+\infty} \frac{(N - \lfloor \Phi_p \rfloor + 1) \cdot \lfloor \Phi_p \rfloor \cdot \lfloor \Phi_p \rfloor^p}{(N - \lfloor \Phi_p \rfloor + 1) + 2 \cdot \sum_{j=1}^{\lfloor \Phi_p \rfloor - 1} c_j} \tag{3.7}$$
$$= \lfloor \Phi_p \rfloor^{p+1} \tag{3.8}$$

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

Proof of claim. Let $i \in [N]$ and $j \in [N-1]$. We prove $d_{e_{i,j}}(f_{e_{i,j}} + 1) \geq \text{cost}(P^1_j)$. To show $d_{e_{i,\lfloor \Phi_p \rfloor + N - j}}(f_{e_{i,\lfloor \Phi_p \rfloor + N - j}} + 1) \geq \text{cost}(P^1_{N-j+1})$ 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

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

where the second equality follows the definition of c_j.

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

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, 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 $\lfloor \Phi_p \rfloor^{p+1} \geq \text{cost}(P_j^1)$. 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

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

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

$$(\lfloor \Phi_p \rfloor - 1)^p \leq \lfloor \Phi_p \rfloor^{p-1} . \tag{3.9}$$

To prove that (3.9) 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 3.7, 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 3.7, $\Phi_p > 1$ 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 (3.9).

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

$$\left(\frac{\ell - 1}{\lfloor \Phi_p \rfloor} \right)^p \cdot \lfloor \Phi_p \rfloor^{p+1} \leq \left(\frac{\ell}{\lfloor \Phi_p \rfloor + 1} \right)^p (\lfloor \Phi_p \rfloor + 1)^p. \tag{3.10}$$

We conclude that (3.10) 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 6

$$\sum_{\ell=1}^{\min\{N,j\}} \text{cost}(P^1_\ell) \leq \sum_{\ell=1}^{j} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \leq w(\bar{P}^1_{s,v_{i,j}}). \tag{3.11}$$

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 $j = i-1 \leq N$, 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

$$w(P^1_{s,v_{i,j}}) = \sum_{\ell=1}^{i-1} \text{cost}(P^1_\ell) = \sum_{\ell=1}^{j} \text{cost}(P^1_\ell). \tag{3.12}$$

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

$$w(P^1_{s,v_{i,j}}) \leq \sum_{\ell=1}^{i} \text{cost}(P^1_\ell) \leq \sum_{\ell=1}^{\min\{N,j\}} \text{cost}(P^1_\ell). \tag{3.13}$$

By combining equations (3.11), (3.12) and (3.13) we obtain that

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

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 6

$$\sum_{\ell=\max\{j-\lfloor\Phi_p\rfloor+2,1\}}^{N} \mathrm{cost}(P^1_\ell) \le \sum_{\ell=j+1}^{N+\lfloor\Phi_p\rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \le w(\bar{P}^1_{s,v_{i,j}}). \tag{3.14}$$

Note that since $v_{i,j}$ belongs to P^1_i, it must be $i-1 \le j \le i+\lfloor\Phi_p\rfloor-1$. 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

$$w(P^1_{v_{i,j},t}) = \sum_{\ell=i+1}^{N} \mathrm{cost}(P^1_\ell) \le \sum_{\ell=\max\{j-\lfloor\Phi_p\rfloor+2,1\}}^{N} \mathrm{cost}(P^1_\ell). \tag{3.15}$$

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

$$w(P^1_{v_{i,j},t}) \le \sum_{\ell=i}^{N} \mathrm{cost}(P^1_\ell) \le \sum_{\ell=\max\{j-\lfloor\Phi_p\rfloor+2,1\}}^{N} \mathrm{cost}(P^1_\ell). \tag{3.16}$$

By combining equations (3.14), (3.15) and (3.16) we obtain that

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

Case(iii): $v_1 \ne s$ and $v_2 \ne t$. 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_i^1, and $h \leq k \leq h + \lfloor \Phi_p \rfloor - 1$. 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 $k + \lfloor \Phi_p \rfloor - 1 \geq h + \lfloor \Phi_p \rfloor - 1 \geq i + \lfloor \Phi_p \rfloor$, 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

$$w(\bar{P}^1_{v_1,v_2}) \geq \sum_{\ell=i+\lfloor \Phi_p \rfloor}^{k+\lfloor \Phi_p \rfloor -1} d_{e_{1,\ell}}(f_{e_{1,\ell}} + 1) \geq \sum_{\ell=i+\lfloor \Phi_p \rfloor}^{h+\lfloor \Phi_p \rfloor -1} d_{e_{1,\ell}}(f_{e_{1,\ell}} + 1). \tag{3.17}$$

Then according to the second inequality of Claim 6, we have

$$\sum_{\ell=i+\lfloor \Phi_p \rfloor}^{h+\lfloor \Phi_p \rfloor -1} d_{e_{1,\ell}}(f_{e_{1,\ell}} + 1) \geq \sum_{\ell=i+1}^{h} w(P^1_\ell). \tag{3.18}$$

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

$$\sum_{\ell=i+1}^{h} \mathrm{cost}(P^1_\ell) \geq w(P^1_{v_1,v_2}) \tag{3.19}$$

By combining inequalities (3.17), (3.18) and (3.19) we have $w(\bar{P}^1_{v_1,v_2}) \geq w(P^1_{v_1,v_2})$.

Subcase(*iii*)*.2:* $i - 1 \leq j \leq i + \lfloor \Phi_p \rfloor - 2$ *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 $h - 1 \geq j - \lfloor \Phi_p \rfloor + 2$, 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

$$w(\bar{P}^1_{v_1,v_2}) \geq \sum_{\ell=j-\lfloor\Phi_p\rfloor+2}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1). \tag{3.20}$$

Then according to the first inequality of Claim 6, we have

$$\sum_{\ell=i}^{h-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h-1} w(P^1_\ell). \tag{3.21}$$

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

$$\sum_{\ell=i}^{h-1} \text{cost}(P^1_\ell) \geq w(P^1_{v_1,v_2}) \tag{3.22}$$

By combining inequalities (3.20), (3.21) and (3.22) 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

$$w(\bar{P}^1_{v_1,v_2}) \geq \sum_{\ell=j-\lfloor\Phi_p\rfloor+2}^{k+\lfloor\Phi_p\rfloor-1} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1). \tag{3.23}$$

Then according to the first inequality of Claim 6, we have

$$\sum_{\ell=i}^{h} d_{e_{1,\ell}}(f_{e_{1,\ell}}+1) \geq \sum_{\ell=i}^{h} w(P^1_\ell). \tag{3.24}$$

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_i^1 and P_h^1. So we can conclude that

$$\sum_{\ell=i}^{h} \text{cost}(P_\ell^1) \geq w(P_{v_1,v_2}^1) \tag{3.25}$$

By combining inequalities (3.23), (3.24) and (3.25) we have $w(\bar{P}_{v_1,v_2}^1) \geq w(P_{v_1,v_2}^1)$. □

3.3 Series-parallel network congestion games with respect to total cost

3.3.1 Upper bound on price of anarchy

In this section, we prove the upper bound on the PoA stated in Theorem 3.2. 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 [25]. 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 $\{C_1, \ldots, C_h\}$ 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 $C_i^- = \{e = (u,v) \in E : (v,u) \in C_i, o_e < f_e\}$.

Recall the parameter $y(\mathcal{D})$ we have defined in Section 3.1. 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, cost$(f)/N$.

Lemma 3.8. *If C_i^+, $i \in [h]$, is an (s,t)-path, then cost$_o(C_i^+) \geq$ 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 3.9. *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*

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

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

$$\sum_{e \in p}(o_e + 1)d_e(o_e + 1) - \sum_{e \in p} o_e d_e(o_e) < \frac{\text{cost}(f)}{N}. \tag{3.26}$$

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

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

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

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

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

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

By summing up these two inequalities we get

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

By Lemma 3.8, since C_i^+ is an (s,t)-path, we have $\text{cost}_o(C_i^+) = \sum_{e \in C_i^+} d_e(o_e) \geq \frac{\text{cost}(f)}{N}$. Thus, by (3.26) we obtain $\text{cost}(o') - \text{cost}(o) < 0$, which contradicts the fact that o is a social optimum. $\square$

By using Lemma 3.9, we can derive a lower bound on $\text{cost}_o(p)$ similar to the lower bound on $\text{cost}_o(C_i^+)$ stated in Lemma 3.8, but with an extra factor of $y(\mathcal{D})$.

Lemma 3.10. *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 $p \in P$,*

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

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

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

where the first inequality follows the definition of $y(\mathcal{D})$ stated in Equation (3.1) and the second inequality follows from Lemma 3.9. □

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

Proof. By Lemma 3.10 we know that given an arbitrary decomposition P of the social optimal flow o, for all $p \in P$, we have $y(\mathcal{D})\text{cost}_o(p) \geq \frac{\text{cost}(f)}{N}$. Then we can conclude that:

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

where the last equality follows from the fact that $|P| = N$. This implies that cost$(f) \leq y(\mathcal{D})$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 [25]. 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 3.12. *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. $\square$

By Proposition 1 in [25], 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 3.13. *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 3.12 $\mathcal{L}$ 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 cost($o'(E_i)$) $<$ 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 cost(o'') $<$ cost(o), contradicting the fact that o is the social optimum. $\square$

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

Lemma 3.14. *If G is series-parallel and E_i is a maximal set in $\mathcal{L}$, then*

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

Proof. According to Lemma 3.13, 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 3.11 we conclude that the lemma holds. $\square$

We are finally ready to prove Theorem 3.2, 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 3.2. 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 cost(f) $\leq y(\mathcal{D})$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 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

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

where the inequality follows from the fact that $y(\mathcal{D}) \geq 1$ and from Lemma 3.14. □

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 3.15. *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 (3.1) we have that for any $x \in \mathbb{N}^+$

$$\begin{aligned}y(\text{Poly-}\,p) &= \sup_{a_0,\ldots,a_p\in\mathbb{R}_{\geq 0},\ x\in\mathbb{N}^+} \frac{(x+1)\sum_{j=0}^{p} a_j(x+1)^j - x\sum_{j=0}^{p} a_j x^j}{\sum_{j=0}^{p} a_j x^j}\\ &= \sup_{a_0,\ldots,a_p\in\mathbb{R}_{\geq 0},\ x\in\mathbb{N}^+} \frac{\sum_{j=0}^{p}\left(a_j\left((x+1)^{j+1} - x^{j+1}\right)\right)}{\sum_{j=0}^{p} a_j x^j}.\end{aligned} \tag{3.27}$$

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 (3.27) by

$$\max_{j\in\{0,\ldots,p\},x\in\mathbb{N}^+} \frac{(x+1)^{j+1} - x^{j+1}}{x^j}. \tag{3.28}$$

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 (3.28) 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 3.2 and Lemma 3.15 we obtain that the PoA of series-parallel network congestion games with polynomial delay functions with highest degree is p is at most $2^{p+1} - 1$.

3.3.2 Lower bound on price of anarchy

In this section, we illustrate how to construct a family of instances that asymptotically achieve the lower bound on the PoA stated in Theorem 3.3. This construction is an extension to polynomial delays of the construction proposed in [30] for affine delays. 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 $i \in [N-1]$. Let $m \in [N-1]$. 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 $j \geq m+1$. We construct a series-parallel (s,t)-network G with delays in $\text{Poly-}\,p$, 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 $e_1, \ldots, e_m$ 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.

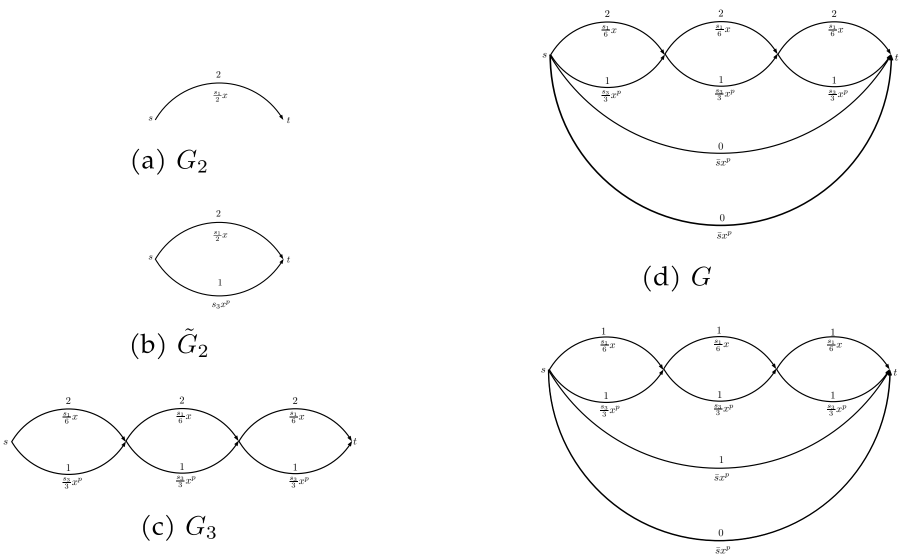

(a) G_2

(b) $\tilde{G}_2$

(c) G_3

(d) G

(e) h

Figure 3.1: 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 $\bar{s} = \frac{s_1+s_2+s_3}{3}$. (3.1(d)) is the output network G and its corresponding PNE flow f. (3.1(a)), (3.1(b)), (3.1(c)) are the intermediate networks and flows according to our construction. (3.1(e)) is the flow h defined in the proof of Theorem 3.3 where $k = 1$.

To prove Theorem 3.3, 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 (3.2) are defined as $r = \frac{m}{N}$, $l = \frac{k}{m}$. Consider the construction described above, which is represented in Fig.3.1. We will first show that this construction satisfies the properties stated in the next two lemmas.

Lemma 3.16. *The (s,t)-flow f has an (s,t)-path $\bar{p}$ with flow value m and cost$_f(\bar{p}) = s_1$.*

Proof. We prove the lemma 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 (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 3.17. *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 cost$(f) =$ 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 3.18. *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:

$$\sum_{j=1}^{N-1} q_j = 1 - q_N = 1 - \frac{1}{2^p(N-1)} \sum_{j=1}^{N-1} q_j. \tag{3.29}$$

By equation (3.29), we have $\frac{2^p(N-1)+1}{2^p(N-1)} \sum_{j=1}^{N-1} q_j \le 1$. This implies that $\sum_{j=1}^{N-1} q_j \le \frac{2^p(N-1)}{2^p(N-1)+1} = \mu(N-1,N)$. 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, $\sum_{j=1}^k q_j \le \mu(k,N)$. Moreover, since $q_k = \frac{1}{2^p(k-1)} \sum_{j=1}^{k-1} q_j$, we have:

$$\sum_{j=1}^{k-1} q_j = \sum_{j=1}^{k} q_j - q_k = \mu(k,N) - \frac{1}{2^p(k-1)} \sum_{j=1}^{k-1} q_j. \tag{3.30}$$

According to (3.30), 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. $\square$

Lemma 3.19. *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.

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

$\square$

We will now use the results stated in the above lemmas to prove Theorem 3.3.

Proof of Theorem 3.3. Consider the network congestion game on the network G defined above. By Lemma 3.16, 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:

$$\begin{aligned}
\text{cost}(f) - \text{cost}(h) &= m\text{cost}_f(\bar{p}) - \left(k\text{cost}_h(\bar{p}) + (m-k)\frac{1}{N}\right) \\
&= s_1 m - \left(\frac{s_1}{m}k^2 + (m-k)\frac{1}{N}\right) && (3.31) \\
&= \left(\frac{s_1}{m}m^2 - \frac{s_1}{m}k^2 - \frac{m-k}{m}ms_1\right) + \frac{m-k}{m}\left(ms_1 - \frac{m}{N}\right) \\
&= \left(\frac{s_1}{m}mk - \frac{s_1}{m}k^2\right) + \frac{m-k}{m}\left(\sum_{i=1}^{m} s_i - \frac{m}{N}\right), && (3.32)
\end{aligned}$$

where equality (3.31) holds since $\sum_{e\in\bar{p}} a_e = \frac{s_1}{m}$. Equality (3.32) holds since the first m of s_i are equal.

By Lemma 3.18 and Lemma 3.19 we have

$$\sum_{i=1}^{m} q_i - \frac{m}{N} = \mu(m,N) - \frac{m}{N} \geq \left[(\sqrt[2p]{r} - r) - \epsilon\right]. \tag{3.33}$$

Now observe that

$$\frac{s_1}{m}m^2 = ms_1 = \frac{m}{N} + (\sum_{1}^{m} s_i - \frac{m}{N}) \geq r + \left[(\sqrt[2p]{r} - r) - \epsilon\right] = (\sqrt[2p]{r} - \epsilon), \tag{3.34}$$

where the inequality follows from (3.33) and the fact that $\sum_1^m s_i = \sum_1^m q_i$.

This implies

$$\frac{s_1}{m}mk - \frac{s_1}{m}k^2 = (l - l^2)\frac{s_1}{m}m^2 \geq (l - l^2)(\sqrt[2p]{r} - \epsilon), \tag{3.35}$$

where the inequality follows from (3.34).

From (3.32) and (3.35)we obtain

$$\begin{aligned}
\text{cost}(f) - \text{cost}(h) &\geq (l - l^2)(\sqrt[2p]{r} - \epsilon) + (1-l)(\sum_{1}^{m} s_i - \frac{m}{N}) \\
&\geq (l - l^2)(\sqrt[2p]{r} - \epsilon) + (1-l)\left[(\sqrt[2p]{r} - r) - \epsilon\right], && (3.36)
\end{aligned}$$

where inequality (3.36) follows from (3.33). By Lemma 3.17 we know that cost(f) $=$ $\sum_1^N s_i = 1$, thus we obtain:

$$\begin{aligned}\text{cost}(h) &\leq 1 - (l - l^2)(\sqrt[2^p]{r} - \epsilon) - (1-l)\left[(\sqrt[2^p]{r} - r) - \epsilon\right] \\ &= 1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r + (1-l^2)\epsilon \qquad (3.37)\end{aligned}$$

To obtain an upper bound on cost(h) we minimize the right-hand-side of (3.37) 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 $\frac{\text{cost}(f)}{\text{cost}(o)} \geq \frac{\text{cost}(f)}{\text{cost}(h)}$, we obtain a lower bound for the PoA. □

We now argue that the worst case PoA is in $\Omega(2^p/p)$. By substituting the expression of l in the denominator of (3.2), we obtain

$$1 + l^2 \sqrt[2^p]{r} - rl - \sqrt[2^p]{r} + r = 1 - \frac{1}{4} r^{2-\frac{1}{2^p}} + r - r^{\frac{1}{2^p}}. \qquad (3.38)$$

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 (3.38) is in $O\left(p/2^p\right)$, which implies that when N goes to infinity the PoA is at least in $\Omega\left(2^p/p\right)$.

3.4 Series-parallel network congestion games with respect to 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.

3.4.1 Upper bound on price of anarchy

We first prove the upper bound on the PoA stated in Theorem 3.4.

Proof of Theorem 3.4. 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 $\max(P_f) \leq \mathrm{cost}_f^+(p_f^i)$ for any $i \in [N]$. Moreover, by (3.3), we have $\mathrm{cost}_f^+(p_f^i) \leq z(\mathcal{D})\mathrm{cost}_f(p_f^i)$. In other words, the most expensive path in P_f has cost no greater than $z(\mathcal{D})$ times the cost of any other path in P_f. Thus we can conclude that

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

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,

$$\begin{aligned} z(\mathcal{D}) \,\mathrm{tot}(P_f) &\leq z(\mathcal{D})y(\mathcal{D}) \,\mathrm{tot}(P_o) && (3.39) \\ &\leq z(\mathcal{D})y(\mathcal{D}) \,\mathrm{tot}(P_{\hat{o}}) && (3.40) \\ &\leq z(\mathcal{D})y(\mathcal{D})(N \cdot \max(P_{\hat{o}})). && (3.41) \end{aligned}$$

Inequality (3.39) directly follows Theorem 3.2. Inequality (3.40) 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 (3.41) holds because $\max(P_{\hat{o}})$ is the maximum player's cost in $P_{\hat{o}}$. $\square$

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

Lemma 3.20. *For the class of polynomial delay functions* $\mathrm{Poly}\text{-}\,p$ *it holds that* $z(\mathrm{Poly}\text{-}\,p) \leq 2^p$.

Proof. By the definition of $z(\mathrm{Poly}\text{-}\,p)$ in (3.3) we have that for any $x \in \mathbb{N}^+$

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

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,

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

$\square$

3.4.2 Lower bound on price of anarchy

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 3.5. 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 cost(P_f)/cost(P_o). Our goal is to construct a new instance on a network G', and to define a PNE $P_{f'}$ and a social optimum $P_{o'}$ with respect to the maximum players' cost, such that

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

We construct G' as follows. First, let $G_1, \ldots, G_N$ 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.

□

3.5 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 chapter 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?

4 PRICE OF ANARCHY IN k-UNIFORM AND PAVING MATROID CONGESTION GAMES

4.1 Introduction

4.1.1 Overview and main results

In this chapter we focus on another combinatorial structure, namely matroids. *Matroid congestion games* are congestion games where each player's strategy set is the set of bases of a given matroid.

Recall that for symmetric k-uniform matroid congestion games with polynomial delays of highest degree p the worst-case PoA is in $O(2^{p(p+1)})$ and in $\Omega(2^p)$ [36]. This indicates that the combinatorial structure of k-uniform matroids significantly limits the inefficiency of equilibria. However, k-uniform matroids are very special matroids, since every subset of the ground set of size at most k is independent. *Are there weaker matroid structures that affect the inefficiency of equilibria?* In this chapter we focus on *paving matroids*, i.e., matroids whose circuits have cardinality greater than or equal to the matroid rank. Unlike k-uniform matroids, paving matroids exhibit a notable predominance within the enumeration of matroids. It has been conjectured that, in an asymptotic sense, the majority of matroids are paving matroids [41]. This conjecture holds if the ground set has size at most 9 [10, 42]. Pendavingh and van der Pol [48] more recently proved that, as the size of the ground set goes to infinity, the ratio of logarithms between the total number of matroids and the number of sparse paving matroids, a subclass of paving matroids, converges to 1.

Our contributions. First, we provide a lower bound of $13/9$ on the worst-case PoA for symmetric paving matroid congestion games with affine delays. This ratio is worse than the previously known best upper bound ≈ 1.41 on the PoA of symmetric congestion games with affine delay functions over k-uniform matroids, which are a subclass of paving matroids. Thus, relaxing the structure of players' strategy sets from uniform matroids to paving matroids can increase the inefficiency of pure Nash equilibria.

Theorem 4.1. *The worst-case PoA of symmetric paving matroid congestion games with affine delay functions is at least* $13/9$*.*

We next turn to the question of finding upper bounds on the PoA of symmetric paving matroid congestion games. Given the class of delay functions $\mathcal{D}$, we define the parameter $z(\mathcal{D})$ as

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

Since the delay functions $d(x)$ are non-negative and non-decreasing, we have $z(\mathcal{D}) \geq 1$. Our first main result is an upper bound on the worst-case PoA in symmetric paving matroid congestion games with delay functions in class $\mathcal{D}$.

Theorem 4.2. *The PoA of symmetric paving matroid congestion games with delay functions in class* $\mathcal{D}$ *is at most* $z(\mathcal{D})^2 \rho(\mathcal{D})$*.*

When $\mathcal{D}$ is the class of polynomial functions of maximum degree p, we have $z(\mathcal{D}) = 2^p$ and $\rho(\mathcal{D}) \in \Theta(p/\ln p)$. Thus, the worst-case PoA is in $O(4^p p/\ln p)$. For $p \geq 6$ our bound is smaller than the worst-case PoA that can be achieved in general symmetric congestion games, that is in $\Theta(p/\ln p)^{p+1}$ [3]. Thus, the worst-case PoA of symmetric congestion games *cannot* be achieved in paving matroids.

We also prove —with a substantially different approach— that this is the case for $p = 1$, i.e., when the delay functions are affine. In this case, the worst-case PoA for general symmetric congestion games is $5/2$.

Theorem 4.3. *The PoA of symmetric paving matroid congestion games with affine delay functions is at most* $17/7$*.*

Finally, the approach used to prove Theorem 4.2 also provides a new upper bound on the worst-case PoA in symmetric k-uniform matroid congestion games with delay functions in class $\mathcal{D}$.

Theorem 4.4. *The PoA of symmetric k-uniform matroid congestion games with delay functions in class $\mathcal{D}$ is at most $z(\mathcal{D})\rho(\mathcal{D})$.*

When $\mathcal{D}$ is the class of polynomial functions of maximum degree p, we obtain that the worst-case PoA is in $O(2^p p/\ln p)$. This significantly improves on the previously known upper bound of $O(2^{p(p+1)})$ [36] and partially closes the gap with the lower bound of $\Omega(2^p)$ [36].

4.1.2 Preliminaries

In this section, we recall some basics of matroid theory and then we introduce some fundamental notions of congestion games.

Matroids. A *matroid* is a pair $(R, \mathcal{I})$ where the *ground set* R consists of a finite set of elements and $\mathcal{I}$ is a nonempty collection of subsets of R such that: (i) if $I \in \mathcal{I}$ and $J \subseteq I$, then $J \in \mathcal{I}$; and (ii) if $I, J \in \mathcal{I}$ and $|I| < |J|$, then $I \cup \{z\} \in \mathcal{I}$ for some $z \in J \setminus I$. Given a matroid $M = (R, \mathcal{I})$, a subset I of R is called *independent* if I belongs to $\mathcal{I}$, and dependent otherwise. A subset $B \subseteq R$ is called a *basis* if B is an inclusion-wise maximal independent subset. That is, $B \in \mathcal{I}$ and there is no $Z \in \mathcal{I}$ with $B \subset Z \subseteq R$. The common size of all bases is called the *rank* of the matroid, denoted by $r(M)$. A *circuit* of a matroid is an inclusion-wise minimal dependent set. For every basis B and every element in $R \setminus B$, there is a unique circuit contained in $B \cup \{x\}$, that is called a *fundamental circuit*.

Next, we introduce the *bijective basis-exchange property*:

Theorem 4.5 ([11])**.** *Let $\mathcal{B}$ be the collection of bases of a matroid. For any $B, B' \in \mathcal{B}$, there is a bijection $\pi : B \to B'$ from B to B', such that for every $x \in B \setminus B'$, $B \setminus \{x\} \cup \{\pi(x)\}$ is a basis.*

A matroid is called *k-uniform* matroid if its independent sets are all the subsets of R of cardinality at most k, i.e. every $k + 1$-element subset of R is a circuit. A matroid is called *paving* matroid if every circuit of M has cardinality $r(M)$ or $r(M) + 1$. The following proposition characterizes paving matroids in terms of their circuits.

Proposition 4.6 ([47]). *Let $\mathcal{C}$ be a collection of non-empty subsets of a set R such that each each member of $\mathcal{C}$ has size either t or $t+1$. Let $\mathcal{C}' \subseteq \mathcal{C}$ consist only of the t-element members of $\mathcal{C}$. Then $\mathcal{C}$ is the set of circuits of a paving matroid on R of rank t if and only if*

1. *if two distinct members C_1 and C_2 of $\mathcal{C}'$ have $t-1$ common elements, then every t-element subset of $C_1 \cup C_2$ is in $\mathcal{C}'$; and*

2. *$\mathcal{C} \setminus \mathcal{C}'$ consists of all the $(t+1)$-element subsets of R that contains no member of $\mathcal{C}'$.*

Matroid congestion games. A *matroid congestion game* is a congestion game where the strategy set of each player i is the set of bases $\mathcal{B}_i$ of a given matroid $M_i = (R_i \subseteq R, \mathcal{I}_i)$. For an arbitrary state s of the matroid congestion game, we denote by B_s^i the strategy of player i in s. A *paving* matroid congestion game is a matroid congestion game where M_i is a paving matroid for all $i \in [N]$. A *k-uniform* matroid congestion game is a congestion game where M_i is a k-uniform matroid for all $i \in [N]$ and $k \in [\min_i |R_i|]$.

4.2 Upper bounds on the PoA for delays in class $\mathcal{D}$

In this section, our goal is to prove Theorems 4.2 and 4.4. For a matroid congestion game over resource set R, we let $G = (R, E)$ be a complete directed graph, where the nodes correspond to the resources in R. Let s and q be two states of the congestion game. We define the following two sets:

$$R^-(s,q) = \{r \in R : s_r > q_r\} \qquad\qquad R^+(s,q) = \{r \in R : s_r < q_r\},$$

and we let $l = \sum_{r \in R^-(s,q)} (s_r - q_r) = \sum_{r \in R^+(s,q)} (q_r - s_r)$. In G, every node $r \in R^-$ has supply $s_r - q_r$, and every node $r \in R^+$ has demand $q_r - s_r$. A (single-commodity) flow $F \in \mathbb{Z}^{R \times R}$ in G is a non-negative vector such that for every node $r \in R$

$$F(\delta^-(r)) - F(\delta^+(r)) = q_r - s_r, \tag{4.1}$$

where $\delta^{-}(r)$ contains all the arcs whose head is r and $\delta^{+}(r)$ contains all the arcs whose tail is r. We call F a (s,q)-*difference flow*. Note that the above definitions can be applied to a generic congestion game. For a matroid congestion game, we can construct a special (s,q)-difference flow F, that we call (s,q)-*exchange flow*, as follows. According to Theorem 4.5, for each pair (B^i_s, B^i_q), there is a bijection $\pi^i(x) : B^i_s \to B^i_q$ such that for every $r \in B^i_s \setminus B^i_q$ there is an unique $\pi^i(r) \in B^i_q \setminus B^i_s$ and $B^i_s \setminus \{r\} \cup \{\pi^i(r)\} \in \mathcal{B}$. Starting from the zero vector, for every $i \in [N]$, $r \in B^i_s \setminus B^i_q$, we add one unit of flow to the arc $(r, \pi^i(r))$ to G in order to obtain F. We observe that F can be decomposed into l paths, each one starting from a node in $R^{-}(s,q)$ and ending at a node in $R^{+}(s,q)$, and carrying one unit of flow. Each path in the exchange flow can be interpreted as a sequence of resource exchanges such that each arc (r, r') in the path corresponds to some player replacing resource r with resource r' in their strategy.

In the next theorem, we consider an (f,o)-exchange flow. For any (u,v)-path from $R^{-}(f,o)$ to $R^{+}(f,o)$, if $\text{cost}_f(u)$ is equal to at least a fraction α of $\text{cost}^{+}_f(v)$, then we can upper bound the ratio between the social costs of f and o by $\alpha\rho(\mathcal{D})$. We recall that the function $\rho(\mathcal{D})$, initially introduced by Roughgarden [52], is defined as $\rho(\mathcal{D}) := \sup_{d \in \mathcal{D}} \rho(d)$, where

$$\rho(d) = \sup_{x \geq y \geq 0} \frac{x d(x)}{y d(y) + (x-y) d(x)}. \tag{4.2}$$

Theorem 4.7. *Let F be an (f,o)-exchange flow. Let $R^{-} = R^{-}(f,o)$ and $R^{+} = R^{+}(f,o)$. For all paths p contained in F from $u \in R^{-}$ to $v \in R^{+}$, if $\alpha \text{cost}^{+}_f(v) \geq \text{cost}_f(u)$ for some $\alpha \geq 1$ then we have $\text{cost}(f) \leq \alpha\rho(\mathcal{D})\text{cost}(o)$.*

Proof. For every resource $r \in R^{-}$, inequality (4.2) and $\alpha \geq 1$ imply

$$\begin{aligned} f_r \text{cost}_f(r) = f_r d_r(f_r) &\leq \rho(\mathcal{D})(o_r d(o_r) + (f_r - o_r) d(f_r)) \\ &\leq \rho(\mathcal{D})(\alpha o_r d(o_r) + (f_r - o_r) d(f_r)) \end{aligned} \tag{4.3}$$

Let $\{p_1, \ldots, p_l\}$ be an arbitrary decomposition of the flow F, where each p_k is from r^{-}_k

to r_k^+ such that $r_k^- \in R^-$ and $r_k^+ \in R^+$. We have

$$\begin{aligned}\sum_{r\in R^+}(o_r - f_r)\text{cost}_o(r) &= \sum_{k=1}^{l}\text{cost}_o(r_k^+) \geq \sum_{k=1}^{l}\text{cost}_f^+(r_k^+)\\ &\geq \sum_{k=1}^{l}\frac{1}{\alpha}\text{cost}_f(r_k^-) = \frac{1}{\alpha}\sum_{r\in R^-}(f_r - o_r)\text{cost}_f(r),\end{aligned} \tag{4.4}$$

where the equalities hold by the definition of F and equality (4.1), the first inequality holds because of the definition of R^+, and the second inequality holds because by assumption. Let $\bar{R} = \{r \in R : f_r = o_r\} = R \setminus (R^- \cup R^+)$.

$$\begin{aligned}\text{cost}(f) &= \sum_{r\in R^-} f_r\text{cost}_f(r) + \sum_{r\in R^+} f_r\text{cost}_f(r) + \sum_{r\in\bar{R}} f_r\text{cost}_f(r)\\ &\leq \sum_{r\in R^-} f_r\text{cost}_f(r) + \sum_{r\in R^+} f_r\text{cost}_o(r) + \sum_{r\in\bar{R}} o_r\text{cost}_o(r)\\ &\leq \rho(\mathcal{D})\sum_{r\in R^-}\alpha o_r\text{cost}_o(r) + \rho(\mathcal{D})\sum_{r\in R^-}(f_r - o_r)\text{cost}_f(r)\\ &\quad + \alpha\rho(\mathcal{D})\sum_{r\in R^+} f_r\text{cost}_o(r) + \sum_{r\in\bar{R}} o_r\text{cost}_o(r)\\ &\leq \rho(\mathcal{D})\sum_{r\in R^-}\alpha o_r\text{cost}_o(r) + \rho(\mathcal{D})\sum_{r\in R^+}(o_r - f_r)\text{cost}_o(r)\\ &\quad + \alpha\rho(\mathcal{D})\sum_{r\in R^+} f_r\text{cost}_o(r) + \sum_{r\in\bar{R}} o_r\text{cost}_o(r)\\ &= \alpha\rho(\mathcal{D})\sum_{r\in R^-\cup R^+} o_r\text{cost}_o(r) + \sum_{r\in\bar{R}} o_r\text{cost}_o(r) \leq \alpha\rho(\mathcal{D})\text{cost}(o).\end{aligned}$$

The first inequality holds because of the definition of R^+ and $\bar{R}$; the second inequality holds because of inequality (4.3) and $\alpha \geq 1$, $\rho(\mathcal{D}) \geq 1$; the third inequality follows by applying (4.4); the last inequality follows because $\alpha \geq 1$, $\rho(\mathcal{D}) \geq 1$. $\square$

We emphasize that the bound on the PoA provided by Theorem 4.7 is not restricted to the class of paving matroids. In fact, the assumption of the theorem involves an exchange flow, which is defined for any matroid, and a parameter α. Thus, for any matroid, if we are able to find such α, we are able to bound the PoA.

The next lemma implies that for k-uniform matroids $\alpha = z(\mathcal{D})$ satisfies the assumption of Theorem 4.7. This lemma is an extension of Lemma 5 in [18] from affine delay functions to general delay functions. Moreover, it can be verified that for polynomial delay functions the bound established in Lemma 4.8 is tight.

Lemma 4.8. *Suppose M is a k-uniform matroid. Let q be an arbitrary state of the game. For every $u \in R^-(f,q)$ and $v \in R^+(f,q)$ we have $z(\mathcal{D})\text{cost}^+_f(v) \geq \text{cost}_f(u)$.*

Proof. Let u^* be the most expensive resource in $R^-(f,q)$, i.e., $\text{cost}_f(r) \leq \text{cost}_f(u^*)$ for every resource $r \in R^-(f,q)$. To prove the lemma, we will show that for every $v \in R^+(f,q)$ we have $z(\mathcal{D})\text{cost}^+_f(v) \geq \text{cost}_f(u^*)$. By contradiction, suppose there exists a resource $v \in R^+(f,q)$ such that

$$z(\mathcal{D})\text{cost}^+_f(v) < \text{cost}_f(u^*). \tag{4.5}$$

Since $q_v > f_v$, we have $f_v < N$, thus there exists at least one player i who does not use v in f, i.e., $v \notin B^i_f$. We claim that, for all $r \in B^i_f$, we have

$$\text{cost}_f(r) \leq \text{cost}^+_f(v). \tag{4.6}$$

This follows from the fact that, since M is a k-uniform matroid $B^i_f \setminus \{r\} \cup \{v\}$ is a basis of M for all $r \in B^i_f$. Thus, if (4.6) did not hold, player i could deviate from $r \in B^i_f$ to v to decrease their cost. As a consequence, $z(\mathcal{D})\text{cost}^+_f(v) < \text{cost}_f(u^*)$ implies that $v \notin B^i_f$. Moreover, recalling that $z(\mathcal{D}) \geq 1$, we have $\text{cost}^+_f(r) \leq z(\mathcal{D})\text{cost}_f(r)$ for all $r \in R$. Combining this with (4.5) and (4.6), we obtain that, for all $r \in B^i_f$

$$\text{cost}^+_f(r) < \text{cost}_f(u^*). \tag{4.7}$$

Note that (4.7) implies $u^* \notin B^i_f$. Since $u^* \in R^-(f,q)$, $f_{u^*} > o_{u^*} \geq 0$, thus there is at least one player j using u^* in f, i.e., $u^* \in B_f^j$. Since M is a k-uniform matroid, $B_f^j \setminus \{u^*\} \cup \{r\}$ is a basis of M for all $r \in B_f^i$. Moreover, since $u^* \notin B_f^i$ and $|B_f^i| = |B_f^j| = k$, we can conclude that $|B_f^i \setminus B_f^j| \geq 1$. I.e. there exists at least one resource $r^* \in B_f^i$ such that $r^* \notin B_f^j$. Thus, by (4.7), player j could deviate from u^* to r^* to decrease their cost. This contradicts the fact that f is a PNE. □

Applying Theorem 4.7 and Lemma 4.8, we can immediately derive Theorem 4.4.

Next, we show that for paving matroids $\alpha = z(\mathcal{D})^2$ satisfies the assumption of Theorem 4.7. To this purpose, we first introduce an auxiliary result.

Lemma 4.9. *Consider a symmetric matroid congestion game with delays in class $\mathcal{D}$. Let f be a PNE, and o a SO. Let v be a resource that is not used by player i in f and let C_v^i be the unique circuit in $B_f^i \cup \{v\}$. Then, for all $r \in C_v^i$ we have $cost_f^+(r) \leq z(\mathcal{D})cost_f(r) \leq z(\mathcal{D})cost_f^+(v)$.*

Proof. Assume that there exists a resource $r \in C_v^i$ such that $\text{cost}_f(r) > \text{cost}_f^+(v)$. Since C_v^i is the unique circuit that satisfies $C_v^i \setminus \{v\} \subseteq B_f^i$, we have that $B_f^i \setminus \{r\} \cup \{v\} \in \mathcal{B}$, i.e., exchanging r and v defines a feasible strategy for player i. By performing this exchange player i is able to lower their cost, thus contradicting the fact that f is a PNE. Thus, we can conclude that for each $r \in C_v^i$ we have $\text{cost}_f(r) \leq \text{cost}_f^+(v)$. This implies that $z(\mathcal{D})\text{cost}_f(r) \leq z(\mathcal{D})\text{cost}_f^+(v)$. Finally, by the definition of $z(\mathcal{D})$, thus we have $\text{cost}_f^+(r) \leq z(\mathcal{D})\text{cost}_f(r)$. □

For an arbitrary state q, consider an (f, q)-exchange flow F and any path contained in it starting from a node $u \in R^-(f, q)$ and ending at a node $v \in R^+(f, q)$. If the matroid is paving, the next lemma implies that $\text{cost}_f^+(v)$ cannot be smaller than a fraction of $\text{cost}_f(u)$.

Lemma 4.10. *Suppose M is a paving matroid with $r(M) = t \geq 1$. Let q be an arbitrary state of the game and let F be an (f, q)-exchange flow. Let $R^- = R^-(f, q)$ and $R^+ = R^+(f, q)$. For all paths p contained in F from $u \in R^-$ to $v \in R^+$, and for every resource r in p we have $cost_f(r) \leq z(\mathcal{D})^2 cost_f^+(v)$.*

Proof. Let r^* be the most expensive resource of path p in f, i.e., $\text{cost}_f(r) \leq \text{cost}_f(r^*)$ for every resource r in p. Since $t \geq 1$ we know that r^* is used by at least one player in f. We will prove $\text{cost}_f(r^*) \leq z(\mathcal{D})^2 \text{cost}_f^+(v)$. By contradiction, suppose

$$\text{cost}_f(r^*) > z(\mathcal{D})^2 \text{cost}_f^+(v). \tag{4.8}$$

Define

$$S = \{r \in R : \text{cost}_f^+(r) < \text{cost}_f(r^*)\}, \quad \bar{S} = \{r \in R : z(\mathcal{D})\text{cost}_f^+(r) < \text{cost}_f(r^*)\}.$$

Since $z(\mathcal{D}) \geq 1$, we have $\bar{S} \subseteq S$. Moreover, we have the following property.

Claim 7. $|\bar{S}| \geq t$.

Proof of claim. Since v is the last node in p, there exists a player j such that $v \notin B_f^j$. Let C_v^j be the fundamental circuit in $B_f^j \cup \{v\}$. By Lemma 4.9, for all $r \in C_v^j$ we have $\text{cost}_f^+(r) \leq z(\mathcal{D})\text{cost}_f^+(v)$. Thus:

$$z(\mathcal{D})\text{cost}_f^+(r) \leq z(\mathcal{D})^2 \text{cost}_f^+(v) < \text{cost}_f(r^*),$$

where the last inequality comes from (4.8). This implies that $C_v^j \subseteq \bar{S}$. Since in a paving matroid of rank t every circuit has size at least t we obtain $|\bar{S}| \geq t$. $\diamond$

Note that $v \in S$, since $z(\mathcal{D}) \geq 1$, and $r^* \notin S$. Since p traverses both r^* and v, there is an arc (a, b) in p such that $a \notin S$ and $b \in S$. Since (a, b) is contained in F there exists a player i such that $a \in B_f^i$, $b \notin B_f^i$ and $B_f^i \setminus \{a\} \cup \{b\} \in \mathcal{B}$.

First, $a \in B_f^i \setminus S = B_f^i \setminus (B_f^i \cap S)$. Thus $1 \leq t - |B_f^i \cap S|$. We have

$$|\bar{S} \setminus B_f^i| = |\bar{S}| - |\bar{S} \cap B_f^i| \geq t - |\bar{S} \cap B_f^i| \geq t - |S \cap B_f^i| \geq t + (1 - t) = 1,$$

where the first inequality follows from Claim 7. Thus $\bar{S} \setminus B_f^i \neq \emptyset$. Let $w \in \bar{S} \setminus B_f^i$. Let C_w^i be the fundamental circuit in $B_f^i \cup \{w\}$. By Lemma 4.9 for all $r \in C_w^i$ we have

$$\text{cost}_f^+(r) \leq z(\mathcal{D})\text{cost}_f^+(w) < \text{cost}_f(r^*),$$

where the last inequality holds because $w \in \bar{S}$.

This implies $C_w^i \subseteq S$. Recall that $C_w^i \setminus \{w\} \subseteq B_f^i$. Since the matroid is paving, $|C_w^i \setminus \{w\}| \geq t-1$. Finally, as $a \in B_f^i \setminus S$ we can conclude that $B_f^i \setminus \{a\} = C_w^i \setminus \{w\} \subseteq S$. Since $b \in S$, we have $B_f^i \setminus \{a\} \cup \{b\} \subseteq S$. We now prove that every t-element subset of S is a circuit. This immediately contradicts the fact that $B_f^i \setminus \{a\} \cup \{b\}$ is a basis.

Claim 8. Every t-element subset of S is a circuit of the paving matroid M.

Proof of claim. Let h be a player such that $r^* \in B_f^h$ and let r be an arbitrary resource in $S \setminus B_f^h$. We show that $B_f^h \setminus \{r^*\} \cup \{r\}$ is a circuit. Consider the fundamental circuit C_r^h in $B_f^h \cup \{r\}$. We argue that r^* is not in C_r^h. If that was the case, we would have $\text{cost}_f^+(r) \geq \text{cost}_f(r^*)$ by Lemma 4.9, which contradicts $r \in S$. Since we have a paving matroid $C_r^h \geq t$, thus $C_r^h = \{r\} \cup B_f^h \setminus \{r^*\}$. This proves that $B_f^h \setminus \{r^*\}$ forms a circuit with every resource $r \in S \setminus B_f^h$. By applying the first statement in Proposition 4.6 we can conclude that every t-element subset of $S \cup B_f^h \setminus \{r^*\}$ is a circuit. By the definition of S we have $r^* \notin S$, so $S \subseteq S \cup B_f^h \setminus \{r^*\}$ and every t-element subset of S is a circuit. $\diamond$

$\square$

Lemma 4.10 implies that for paving matroids $\alpha = z(D)^2$ satisfies the assumption of Theorem 4.7. Thus, Theorem 4.2 directly follows.

Remark 4.11. It can be verified that the bound of Lemma 4.10 is tight for polynomial delay functions, however we conjecture that the bound of Theorem 4.2 is *not* tight for the same class of delays. In fact, instances where the bound of Lemma 4.10 is tight can have PoA smaller than the upper bound of Theorem 4.2. An intuitive explanation is the following: when the bound in Lemma 4.10 is tight, in the PNE there is an "expensive" resource used by many players and a "cheap" resource used by few players. For this state to be a PNE, the circuits of the matroid must prevent single player deviations where the expensive resource is replaced by the cheap one. The existence of these circuits requires the existence of other resources with comparable costs both in the PNE and in the SO (this is implied by Lemma 4.9). As a result, the PoA in these instances will be lower than the upper bound of Theorem 4.2.

4.3 Lower bound on the PoA of paving matroid congestion games with affine delays

In this section, we consider symmetric paving matroid congestion game with affine delays, i.e., we assume that the delay function of each resource $r \in R$ is of the form $d_r(x) = a_r x + b_r$ with $a_r \geq 0$ and $b_r \geq 0$. Our goal is to prove Theorem 4.1, stating that the worst-case PoA is at least $13/9$. This lower bound is higher than the previously best known lower bound of ≈ 1.35, which is achieved in the symmetric k-uniform matroid congestion games [18]. Moreover, this lower bound indicates that the upper bound of ≈ 1.41 for symmetric k-uniform matroid congestion games does not hold for paving matroids.

Proof. [Proof of Theorem 4.1] We prove the theorem by constructing an instance of a symmetric paving matroid congestion game with affine delays that achieves the PoA of 13/9. Let $R = \{r_1\} \cup R_2 \cup R_3$, where $R_2 = \{r_2, r_3, r_4, r_5\}$ and $R_3 = \{r_6, \ldots, r_{13}\}$. Let

$$C_1 = \{\{r_1, r_{6+2i}, r_{6+2i+1}\} : \forall i \in \{0, 1, 2, 3\}\},$$

$$C_2 = \{S \subset R : |S| = 4 \text{ and } S' \not\subset S, \forall S' \in C_1\}.$$

Let $C = C_1 \cup C_2$. Using Proposition 4.6 with $\mathcal{C}' = C_1$ and $\mathcal{C} = C$ we can easily check that C is the set of circuits for a paving matroid of rank 3 defined over R.

Next we define a symmetric congestion game over $\mathcal{M}$. Let the delay function of r_1 be $d_{r_1}(x) = 1$, and for $i \in \{2, 3, \ldots, 13\}$ let $d_{r_i}(x) = x$. Let the number of players be $N = 6$. The strategy set of each player is the set of bases of the paving matroid. In a PNE, players 1 and 2 select resources $\{r_1, r_2, r_3\}$ and for $i \in \{3, 4, 5, 6\}$, player i selects resources $\{r_4, r_6, r_7\}$, $\{r_4, r_8, r_9\}$, $\{r_5, r_{10}, r_{11}\}$, $\{r_5, r_{12}, r_{13}\}$, respectively. Note that players will not deviate from r_4 or r_5 to r_1, since this would form a circuit in C_1. The social cost of this PNE state is 26. In the SO, each player $i \in [N]$ selects resources $\{r_1, r_{1+i}, r_{7+i}\}$. It can be easily checked that those strategies contain no circuit and the social cost is 18. Thus, the PoA of this instance is at least $26/18 = 13/9$. □

4.4 Upper bound on the PoA of paving matroid congestion games with affine delays

In this section, we prove Theorem 4.3. Consider a symmetric matroid congestion game with N players over resource set R, and suppose that every delay function is affine. Let s and q be two arbitrary states of the game such that cost$(s) \geq$ cost(q), and let $R^- = R^-(s, q)$, $R^+ = R^+(s, q)$. We consider the graph G defined in Section 4.2, where each node $r \in R^-$ has supply $s_r - q_r$ and each node $r \in R^+$ has demand $q_r - s_r$, and we let Φ be a (s, q)-difference flow in G. The following theorem identifies some special properties of Φ that can be used to upper bound $\frac{\text{cost}(s)}{\text{cost}(q)}$. The proof of the theorem is deferred to the end of the section.

Theorem 4.12. *Suppose that Φ is an acyclic (s, q)-exchange flow satisfying the following properties:*

1. *For every arc (u, v) with positive flow in Φ, $cost_s(u) \leq cost_s^+(v)$.*

2. *For every path p from $u \in R^-$ to $v \in R^+$, $cost_s^+(v) \geq \frac{1}{4} cost_s(u)$.*

3. *Let (v, w) be an arc with positive flow in Φ. If for every path to v starting at a node $u \in R^-$ we have $cost_s(v) \geq \frac{1}{2}cost_s(u)$, then $w \notin R^+$.*

4. *For all $r \in R^+$, $s_r = 0$ and $\Phi(\delta^+(r)) = 0$.*

5. *For all $r \notin R^+$, the delay function of r is linear.*

Then $cost(s)/cost(q) \leq 17/7$.

Now consider a symmetric paving matroid congestion game with N players over resource set R, and suppose that the delay functions $d = (d_r)_{r \in R}$ are affine. Let f and o be a PNE and a SO, respectively, that achieve the PoA. We consider an (f, o)-exchange flow F. We then apply five steps, to map $\S = (R, d, f, o, F)$ to a tuple $\S' = (R', d', s, q, \Phi)$ that defines a symmetric 1-uniform matroid congestion game over R' with affine delays $d' = (d'_r)_{r \in R}$, where s and q are two states of the game, and Φ is a (s, q)-exchange flow satisfying the assumptions in Theorem 4.12, and such that

$$\frac{\text{cost}(f)}{\text{cost}(o)} \leq \frac{\text{cost}(s)}{\text{cost}(q)}.$$

Then using Theorem 4.12 we can conclude that the worst-case PoA of symmetric paving matroid congestion games is at most 17/7.

Let $\S^0 = (R, d, f, o, F)$. F is an (f, o)-exchange flow of a matroid congestion game, thus for every arc (u, v) with positive flow in F there exists a player i who could replace resource u with resource v in their strategy. Since f is a PNE, player i is not able to decrease their cost by exchanging u and v, implying that F satisfies property 1. Moreover, since for affine delays $z(\mathcal{D}) = 2$, Lemma 4.10 implies that also property 2 is satisfied. We apply the following four steps, that preserve properties 1 and 2. Moreover, the construction guarantees $\sum_{r \in R} s_r = \sum_{r \in R} q_r$ in every step. This implies that in every step we can construct an instance of a symmetric 1-uniform matroid congestion game on resource set R where s and q are two states that are obtained by assigning players to resources so that for each

$r \in R$ we have s_r players using r in s and q_r players using r in q. The corresponding (s,q)-exchange flow is redefined accordingly. Note that s and q are not necessarily a PNE and a SO of the game.

Step 1. First, we let $s = f$, $q = o$ and $\Phi = F$. We redefine (R, d, s, q, Φ) as follows. For every resource $v \in R^+(f, o)$ such that $f_v > 0$, we add a new resource v' with constant delay equal to $\text{cost}_o(v)$. We set $s_v = q_v = f_v$, $s_{v'} = 0$ and $q_{v'} = o_v - f_v > 0$. Note that $q_{v'} > s_{v'}$, i.e., $v' \in R^+(s,q)$, while $q_v = s_v$, i.e., $v \notin R^+(s,q)$. Moreover we define the flow Φ on arc (v, v') to be $o_v - f_v$. At the end, Φ is a (s,q)-exchange flow that satisfies property 4.

Finally we show that $\frac{\text{cost}(f)}{\text{cost}(o)} \leq \frac{\text{cost}(s)}{\text{cost}(q)}$ after Step 1. Denote the set of nodes we added in this step by V'. According to the construction in Step 1, we have

$$\text{cost}(s) = \sum_{r \in R} s_r \text{cost}_s(r) + \sum_{r \in V'} s_r \text{cost}_s(r) = \sum_{r \in R} f_r \text{cost}_f(r) + 0 = \text{cost}(f),$$

and

$$\begin{aligned}
\text{cost}(q) &= \sum_{r \in R \setminus V} q_r \text{cost}_q(r) + \sum_{r \in V} q_r \text{cost}_q(r) + \sum_{r \in V'} q_r \text{cost}_q(r) \\
&= \sum_{r \in R \setminus V} o_r \text{cost}_o(r) + \sum_{r \in V} f_r \text{cost}_q(r) + \sum_{r \in V} (o_r - f_r) \text{cost}_o(r) \\
&< \sum_{r \in R \setminus V} o_r \text{cost}_o(r) + \sum_{r \in V} f_r \text{cost}_o(r) + \sum_{r \in V} (o_r - f_r) \text{cost}_o(r) = \text{cost}(o),
\end{aligned}$$

where the inequality holds because $\text{cost}_q(r) = d_r(f_r) < d_r(o_r) = \text{cost}_o(r)$. By combining the above inequalities we obtain $\frac{\text{cost}(f)}{\text{cost}(o)} \leq \frac{\text{cost}(s)}{\text{cost}(q)}$.

Step 2. For each resource $v \in R^+(s,q)$ receiving $t_1, \ldots, t_h$ units of flow from $h \geq 2$ resources $u_1, \ldots, u_h$ through arcs $(u_1, v), \ldots, (u_h, v)$ in Φ, we redefine (R, d) by replacing v with h new nodes $v_1, \ldots, v_h$, each having delay function d_v. We redefine (s, q) by setting $s_{v_i} = 0$ and $q_{v_i} = t_i$ for all $i \in [h]$. Next, we redefine Φ by replacing arc (u_i, v) with (u_i, v_i) having flow value t_i, for all $i \in [h]$. After this step, for each $v \in R^+(s,q)$ there is only one resource sending flow to v.

Let $(s, q), (s', q')$ denote the input and output states of Step 2, respectively. We show that $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$ holds after Step 2. For each $v \in R^+(s, q)$ that we selected in Step 2, we replaced it with $v_1, \ldots, v_h$. By the construction we have:

$$s_v \text{cost}_s(v) = \sum_{i=1}^{h} s'_{v_i} \text{cost}_{s'}(v_i) = 0,$$

and

$$q_v \text{cost}_q(v) = \sum_{i=1}^{h} q'_{v_i} \text{cost}_q(v) = \sum_{i=1}^{h} q'_{v_i} d_v(q_v) \geq \sum_{i=1}^{h} q'_{v_i} d_v(q'_v) = \sum_{i=1}^{h} q'_{v_i} \text{cost}_{q'}(v_i).$$

Thus, the social cost of s stays the same and the social cost of q decreases after Step 2, so we have $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$.

Step 3. For each resource $v \in R^+(s, q)$, let r^* be the most expensive resource in $R^-(s, q)$ that is connected to v along a path carrying at least one unit of flow in Φ. Let u be the only resource sending flow to v in Φ, and let h be the flow of Φ on arc (u, v). If $\text{cost}^+_s(v) > \frac{1}{2}\text{cost}_s(r^*)$, we redefine (R, d) by replacing v with h new nodes $v_1, \ldots, v_h$ having delay function $\frac{1}{2}\text{cost}^+_s(v)x$ for $i \in [h]$. Moreover, we add h new resource $w_1, \ldots, w_h$ with constant delay function $\frac{1}{2}\text{cost}^+_s(v)$ for $i \in [h]$. We redefine (s, q) by setting $s_{v_i} = 1$, $s_{w_i} = 0$ and $q_{v_i} = q_{w_i} = 1$ for $i \in [h]$. Thus, property 4 is preserved. Finally, we redefine Φ by setting to one the flow of arcs (u, v_i) and (v_i, w) for $i \in [h]$. We repeat this step until for all $v \in R^+(s, q)$ we have $\text{cost}^+_s(v) \leq \frac{1}{2}\text{cost}_s(r^*)$, thus achieving property 3.

As in Step 2, let $(s, q), (s', q')$ denote the input and output states of each iteration in Step 3, respectively. We show that $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$ holds after each iteration of Step 3. Note that for each $v \in R^+(s, q)$ that we selected in an iteration of Step 3, v is replaced by $v_1, \ldots, v_h$

and $w_1, \ldots, w_h$. By our construction we have:

$$s_v \text{cost}_s(v) = 0 < \sum_{i=1}^{h} (v_i \text{cost}_{s'}(v_i) + w_i \text{cost}_{s'}(w_i)) = \sum_{i=1}^{h} \frac{1}{2} \text{cost}_s^+(v),$$

and

$$\begin{aligned} q_v \text{cost}_q(v) = h\text{cost}_q(v) \geq h\text{cost}_s^+(v) = \sum_{i=1}^{h} \frac{1}{2} \text{cost}_s^+(v) + \sum_{i=1}^{h} \frac{1}{2} \text{cost}_s^+(v) \\ = \sum_{i=1}^{h} q'_{v_i} \text{cost}_{q'}(v_i) + \sum_{i=1}^{h} q'_{w_i} \text{cost}_{q'}(w_i). \end{aligned}$$

The above inequalities imply that after each iteration of Step 3 the social cost of s increases and the social cost of q decreases, so we have $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$.

Step 4. For every resource $r \notin R^+(s, q)$, suppose $d_r(x) = ax + b$ where $a, b \geq 0$. We redefine the delay function of r as $\frac{\text{cost}_s(r)}{s_r} x = \frac{as_r + b}{s_r} x$.

Next we show that $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$, where $(s, q), (s', q')$ are the input and output states of Step 4, respectively. According to the definition of the new delay functions, it is easy to conclude that $\text{cost}(s) = \text{cost}(s')$. For every resource $r \in R \setminus R^+(s', q')$, since we have $s_r = s'_r \geq q'_r = q_r$, then $\text{cost}_{q'}(r) \leq \text{cost}_q(r)$. For every resource $r \in R^+(s', q')$, since we did not change the associated delay function, we have $\text{cost}_{q'}(r) = \text{cost}_q(r)$. Thus, we can conclude that $\text{cost}(q') \leq \text{cost}(q)$, implying $\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\text{cost}(s')}{\text{cost}(q')}$.

Step 5. We delete all the cycles in Φ to make the flow acyclic. At the end, we set $\S' = \{R, d, s, q, \Phi\}$ and $\S = \S^0$. Thus, we achieve property 5.

Based on our discussion we obtain the following lemma.

Lemma 4.13. *$\S'$ satisfies the six assumptions in Theorem 4.12 and $\frac{cost(f)}{cost(o)} \leq \frac{cost(s)}{cost(q)}$.*

Remark 4.14. The construction that we use in the proof of Theorem 4.3 relies on Lemma 4.10 to satisfy property 2 in Theorem 4.12. As discussed in Remark 4.11, although there exist instances where the bound in Lemma 4.10 is tight, these instances might still have

PoA smaller than the upper bound of Theorem 4.3. Thus, we conjecture that the upper bound of Theorem 4.3 is not tight.

We are now left with proving Theorem 4.12.

Proof. [Proof of Theorem 4.12]By property 5, every node $r \notin R^+$ has a linear delay function. Thus, for all $r \notin R^+$ we have

$$\text{cost}_s(r) = \frac{s_r}{s_r + 1}\text{cost}_s^+(r), \tag{4.9}$$

$$\text{cost}_q(r) = \frac{q_r}{s_r + 1}\text{cost}_s^+(r). \tag{4.10}$$

Then we can write

$$1 \leq \frac{\text{cost}(s)}{\text{cost}(q)} = \frac{\sum_{r \in R \setminus R^+} s_r \text{cost}_s(r)}{\sum_{r \in R} q_r \text{cost}_q(r)} \tag{4.11}$$

$$= \frac{\sum_{r \in R \setminus R^+} \frac{s_r^2}{s_r+1}\text{cost}_s^+(r)}{\sum_{r \in R \setminus R^+} \frac{q_r^2}{s_r+1}\text{cost}_s^+(r) + \sum_{r \in R^+} q_r d_r(q_r)}, \tag{4.12}$$

$$\leq \frac{\sum_{r \in R \setminus R^+} \frac{s_r^2}{s_r+1}\text{cost}_s^+(r)}{\sum_{r \in R \setminus R^+} \frac{q_r^2}{s_r+1}\text{cost}_s^+(r) + \sum_{r \in R^+} q_r \text{cost}_s^+(r)}. \tag{4.13}$$

Note that equality (4.11) holds because we have $s_r = 0$ for all $r \in R^+$ according to property 4, while equality (4.12) is implied by (4.9) and (4.10) and inequality (4.13) follows from the fact that $d_r(q_r) \geq \text{cost}_s^+(r)$ for all $r \in R^+$.

Let $r \in R$. We define $\lambda(r) = \text{cost}_s(r) - \frac{1}{2}\text{cost}_s^+(r)$. Let $p = r_0, \ldots, r_k$ be a path in Φ carrying one unit of flow, where $r_0 \in R^-$ and $r_k \in R^+$. For each $h \in [k-1]$ and $i \in [h-1]$ we define:

$$\Psi^0(p,h) := \left(\frac{1}{2}\right)^h \text{cost}_s(r_0), \qquad \Psi^i(p,h) := \left(\frac{1}{2}\right)^i \lambda(r_{h-i}),$$

$$\Omega(p,0) := \sum_{j=1}^{k-1} \Psi^0(p,j), \qquad \Omega(p,i) := \sum_{j=i+1}^{k-1} \Psi^{j-i}(p,j) = \sum_{j=i+1}^{k-1} \left(\frac{1}{2}\right)^{j-i} \lambda(r_i).$$

Moreover, we set $\Omega(p, k-1) = 0$, $\Psi(p, 0) = 0$, and for $h \in [k-1]$ we let $\Psi(p, h) := \sum_{i=0}^{h-1} \Psi^i(p, h)$. It can be checked that

$$\sum_{j=0}^{k-1} \Psi(p, j) = \sum_{j=0}^{k-1} \Omega(p, j). \tag{4.14}$$

Claim 9. We have that $\Psi(p, h) \leq \frac{1}{2}\text{cost}_s^+(r_h)$.

Proof of claim. We prove the claim by induction on h. Let $h = 1$. Since (r_0, r_1) is an arc in the path p, by property 1 we have

$$\Psi(p, 1) = \frac{1}{2}\text{cost}_s(r_0) \leq \frac{1}{2}\text{cost}_s^+(r_1).$$

Now assume that for r_h with $h < k-1$, the claim holds. Then for r_{h+1}, we have

$$\begin{aligned}
\frac{1}{2}\text{cost}_s^+(r_{h+1}) &\geq \frac{1}{2}\text{cost}_s(r_h) && (4.15)\\
&= \frac{1}{2}\lambda(r_h) + \frac{1}{4}\text{cost}_s^+(r_h) && (4.16)\\
&\geq \frac{1}{2}\lambda(r_h) + \frac{1}{2}\Psi(p, h) && (4.17)\\
&= \frac{1}{2}\lambda(r_h) + \frac{1}{2}\left(\left(\frac{1}{2}\right)^h \text{cost}_s(r_0) + \sum_{i=1}^{h-1}\left(\frac{1}{2}\right)^i \lambda(r_{h-i})\right)\\
&= \left(\frac{1}{2}\right)^{h+1} \text{cost}_s(r_0) + \sum_{i=1}^{h}\left(\frac{1}{2}\right)^i \lambda(r_{h+1-i}) = \Psi(p, h+1),
\end{aligned}$$

where (4.15) follows from applying property 1 to the arc (r_h, r_{h+1}) in the path p, equality (4.16) follows from the definition of $\lambda(r_h)$, and inequality (4.17) holds because of our inductive hypothesis. $\diamond$

Now let $P = \{p_1, \ldots, p_l\}$ be an arbitrary decomposition of the flow Φ where each path starts at a node in R^- and ends at a node in R^+ and carries one unit of flow. By property 4, $\Phi(\delta^+(r)) = 0$ for each $r \in R^+$, thus in every path $p \in P$ the only node in R^+ is the sink of the path, denoted by $t(p)$. Moreover, for each resource $r \in R$ we denote by $P(r)$ the paths in P that contain r and by $P^0(r)$ the paths in P starting at r. Finally, for each resource $r \in R$ and path $p \in P(r)$ we use the notation $p(r)$ to identify the position of r in p, precisely $p(r) = 0$ if r is the start node of p, and $p(r) = i$ if r is the i-th node appearing after the start node of p. After summing up inequalities (4.14) for every path $p \in P$ we have

$$\sum_{r \in R \setminus R^+} \sum_{p \in P(r)} \Psi(p, p(r)) = \sum_{r \in R \setminus R^+} \sum_{p \in P(r)} \Omega(p, p(r)). \tag{4.18}$$

For each resource $r \in R$, we define $a_r := \Phi(\delta^-(r))$, $b_r := \Phi(\delta^+(r)) - \Phi(\delta^-(r))$. Note that we have $a_r = |P(r) \setminus P^0(r)|$ and $b_r = |P^0(r)|$ because Φ is acyclic and thus each path $p \in P$ is simple. For each $r \in R \setminus R^+$ and each path in $P(r) \setminus P^0(r)$ we apply Claim 9. Summing up we obtain

$$\Theta_r := \frac{a_r}{2}\text{cost}_s^+(r) - \sum_{p \in P(r) \setminus P^0(r)} \Psi(p, p(r)) \geq 0.$$

Since the fraction in (4.13) is at least 1, by subtracting the non-negative constant $\sum_{r \in R \setminus R^+} \Theta_r$ to both the numerator and the denominator we obtain an upper bound. By using (4.18) and the fact that $\Psi(p, 0) = 0$ for every $p \in P$ we obtain

$$\frac{\text{cost}(s)}{\text{cost}(q)} \leq \frac{\sum_{r \in R \setminus R^+} A_r}{\sum_{r \in R \setminus R^+} B_r} \leq \max_{r \in R \setminus R^+} \frac{A_r}{B_r}, \tag{4.19}$$

where

$$A_r = \frac{(s_r)^2}{s_r + 1}\text{cost}_s^+(r) + \sum_{p \in P(r)} \Omega(p, p(r)) - \frac{a_r}{2}\text{cost}_s^+(r), \tag{4.20}$$

$$B_r = \frac{(q_r)^2}{s_r + 1}\text{cost}_s^+(r) + \sum_{p \in P(r)} \Omega(p, p(r)) + \sum_{p \in P^0(r)} \text{cost}_s^+(t(p)) - \frac{a_r}{2}\text{cost}_s^+(r). \tag{4.21}$$

In the remaining part of the proof we will show that $\frac{A_r}{B_r} \leq \frac{17}{7}$ for all $r \in R \setminus R^+$.

Let $z_r := s_r - a_r - b_r = q_r - a_r$. Thus we have $s_r = z_r + a_r + b_r$ and $q_r = z_r + a_r$. First, we remark that $z_r \geq 0$ for all $r \in R$. This is because Φ is a (s, q)-exchange flow, thus for every node $r \in R$ we have that $\Phi(\delta^+(r))$ is exactly the number of players using r in s and not in q. This number is clearly upper bounded by s_r, the number of players using r in s, thus $s_r \geq \Phi(\delta^+(r)) = a_r + b_r$, implying $z_r \geq 0$. For $r \in R \setminus R^+$ we have $s_r \geq q_r$. If $q_r = s_r$ then we have $B_r - A_r = \sum_{p \in P^0(r)} \text{cost}_s^+(t(p)) \geq 0$, which implies that $\frac{A_r}{B_r} \leq 1$. Thus, to upper bound $\frac{A_r}{B_r}$ we now assume that $s_r \geq q_r + 1$, i.e., $r \in R^-$. Since $s_r = z_r + a_r + b_r$ and $q_r = z_r + a_r$, this implies that we have $b_r \geq 1$. Moreover, since $s_r \geq 1$ we have that $\lambda(r) = \text{cost}_s(r) - \frac{1}{2}\text{cost}_s^+(r) = (\frac{s_r}{s_r+1} - \frac{1}{2})\text{cost}_s^+(r) \geq 0$. This implies:

$$\Omega(p, p(r)) \geq 0 \qquad \forall\, p \in P(r) \setminus P^0(r) \tag{4.22}$$

$$\Omega(p, p(r)) \geq \frac{1}{2}\lambda(r) \qquad \forall\, p \in P(r) \setminus P^0(r) : t(p) > p(r) + 1. \tag{4.23}$$

Let

$$\begin{aligned}
P_1^0(r) :=& \{p \in P^0(r) : \text{cost}_s^+(t(p)) \geq \text{cost}_s(r)\} \\
P_2^0(r) :=& \{p \in P^0(r) : \frac{1}{2}\text{cost}_s(r) \leq \text{cost}_s^+(t(p)) < \text{cost}_s(r)\} \\
P_3^0(r) :=& \{p \in P^0(r) : \frac{1}{4}\text{cost}_s(r) \leq \text{cost}_s^+(t(p)) < \frac{1}{2}\text{cost}_s(r)\}.
\end{aligned}$$

Note that we have $P^0 = P_1^0 \cup P_2^0 \cup P_3^0$. In fact, Lemma 4.10 and $z(\mathcal{D}) = 2$ for the class of affine delay functions imply $\text{cost}_s^+(r) \leq 4\text{cost}_s^+(t(p))$ for every path $p \in P^0$.

First, recalling the definition of $\Omega(p, 0)$ and the fact that delay functions are nonnegative, we obtain that for every path $p \in P^0(r)$

$$\Omega(p, 0) \geq 0. \tag{4.24}$$

Secondly, we prove that for every path $p \in P_2^0(r)$ we have

$$\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r) \geq 0. \tag{4.25}$$

In fact, for every path $p \in P_2^0(r)$, there must exists at least one resource between r and $t(p)$. Otherwise $(r, t(p))$ would be an arc in p and $\text{cost}_s^+(t(p)) < \text{cost}_s(r)$, which contradicts property 1. Thus, from the definition of $\Omega(p,0)$ and the fact that p has at least three nodes, we have $\Omega(p,0) \geq \frac{1}{2}\text{cost}_s(r)$, which implies (4.25).

Finally, for every path $p \in P_3^0(r)$ there must exists at least two resources between r and $t(p)$. Otherwise, if there is no resource between them, then $(r, t(p))$ would be an arc in p and $\text{cost}_s^+(t(p)) < \text{cost}_s(r)$, which contradicts property 1. If there is one resource r' between r and $t(p)$, then we have two edges (r, r') and $(r', t(p))$. By property 1 and the definition of $z(\mathcal{D})$ we have $\text{cost}_s(r) \leq \text{cost}_s^+(r') \leq z(\mathcal{D})\text{cost}_s(r')$. Since for affine delays $z(\mathcal{D}) = 2$, so $\frac{1}{2}\text{cost}_s(r) \leq \text{cost}_s(r')$. Because $p \in P_3^0(r)$, we also have $\text{cost}_s^+(t(p)) < \frac{1}{2}\text{cost}_s(r) \leq \text{cost}_s(r')$, which contradicts property 1 on the edge $(r', t(p))$. Thus, we have $\Omega(p,0) \geq (\frac{1}{2} + \frac{1}{4})\text{cost}_s(r)$ which implies

$$\Omega(p,0) - \frac{3}{4}\text{cost}_s(r) \geq 0. \tag{4.26}$$

From now on we denote by r^* the most expensive resource with respect to state s among all the resources $u \in R^-$ such that there exists a path from u to r in the directed graph induced by Φ. We denote by p^* a path from r^* to r in this graph. Next we need to analyze the following two cases.

Case(i): cost$_s(r) \leq \frac{1}{2}$**cost**$_s(r^*)$

In this case, we show that if $\frac{A_r}{B_r} \geq 1$, then $\frac{A_r}{B_r} \leq \frac{7}{3}$. We first argue that for every path $p \in P^0(r)$, we have

$$\text{cost}_s^+(t(p)) \geq \frac{1}{2}\text{cost}_s(r). \tag{4.27}$$

By contradiction, suppose there exists a path $p \in P^0(r)$ such that $\text{cost}_s^+(t(p)) < \frac{1}{2}\text{cost}_s(r)$. Then, since we are assuming $\text{cost}_s(r) \leq \frac{1}{2}\text{cost}_s(r^*)$, we have that $\text{cost}_s^+(t(p)) < \frac{1}{4}\text{cost}_s(r^*)$. Then we can combine the path p^* from r^* to r and the path p from r to $t(p)$ to obtain a path from $r^* \in R^-$ to $t(p) \in R^+$ that carries at least one unit of flow in Φ. This contradicts to the property 2.

By (4.22) we have $\sum_{p \in P(r) \setminus P^0(r)} \Omega(p, p(r)) \geq 0$. Since $A_r/B_r \geq 1$ we can subtract from both the numerator and the denominator this nonnegative constant and obtain an upper bound. Thus, by using (4.20) and (4.21) we get

$$\frac{A_r}{B_r} \leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} \Omega(p, 0) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} (\Omega(p, 0) + \text{cost}_s^+(t(p))) - \frac{a_r}{2}\text{cost}_s^+(r)}. \tag{4.28}$$

Because of inequality (4.27), we have $P^0(r) = P_1^0(r) \cup P_2^0(r)$. Thus, from (4.28) we obtain

$$\frac{A_r}{B_r} \leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{i=1}^2 \sum_{p \in P_i^0(r)} \Omega(p, 0) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{i=1}^2 \sum_{p \in P_i^0(r)} \left(\Omega(p, 0) + \frac{1}{2^{i-1}}\text{cost}_s(r)\right) - \frac{a_r}{2}\text{cost}_s^+(r)}. \tag{4.29}$$

To upper bound the right-hand-side of (4.29) we do the following. First, for every $p \in P_1^0(r)$, we subtract $\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r)$ from the numerator and subtract $\Omega(p, 0)$ from the denominator. Because $\Omega(p, 0) \geq 0$ for all $p \in P_1^0(r)$ and $\frac{A_r}{B_r} \geq 1$, this will increase the right-hand-side of inequality (4.29). Secondly, for every $p \in P_2^0(r)$, we subtract $\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r)$ from both the numerator and the denominator. This will also increase the right-hand-side of (4.29) because $\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r) \geq 0$ and $\frac{A_r}{B_r} \geq 1$. We obtain

$$\frac{A_r}{B_r} \leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} \frac{1}{2}\text{cost}_s(r) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} \text{cost}_s(r) - \frac{a_r}{2}\text{cost}_s^+(r)} \tag{4.30}$$

$$= \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} \frac{1}{2}\frac{(s_r)}{s_r+1}\text{cost}_s^+(r) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p \in P^0(r)} \frac{(s_r)}{s_r+1}\text{cost}_s^+(r) - \frac{a_r}{2}\text{cost}_s^+(r)}, \tag{4.31}$$

where equation (4.31) follows from (4.9) and $r \notin R^+$. Finally, we rewrite the right-hand-side of (4.31) by factoring out $\text{cost}_s^+(r)$ and exploiting $|P^0(r)| = b_r$, $s_r = a_r + b_r + z_r$, $q_r = a_r + z_r$. Next, we derive an upper bound by considering the maximum over all possible values of a_r, b_r and z_r. We obtain

$$\frac{A_r}{B_r} \leq \frac{\frac{(a_r+b_r+z_r)^2}{a_r+b_r+z_r+1} + b_r\frac{1}{2}\frac{(a_r+b_r+z_r)}{a_r+b_r+z_r+1} - \frac{a_r}{2}}{\frac{(a_r+z_r)^2}{a_r+b_r+z_r+1} + b_r\frac{(a_r+b_r+z_r)}{a_r+b_r+z_r+1} - \frac{a_r}{2}} \leq \max_{a,z\geq 0, b\geq 1} \frac{\frac{(a+b+z)^2}{a+b+z+1} + \frac{1}{2}b\frac{(a+b+z)}{a+b+z+1} - \frac{a}{2}}{\frac{(a+z)^2}{a+b+z+1} + b\frac{(a+b+z)}{a+b+z+1} - \frac{a}{2}} \leq \frac{7}{3}.$$

Case(ii): $\mathbf{cost}_s(r) > \frac{1}{2}\mathbf{cost}_s(r^*)$ In this case, we show that if $\frac{A_r}{B_r} \geq 1$, then $\frac{A_r}{B_r} \leq \frac{17}{7}$. Let p be a path in $P(r) \setminus P^0(r)$. We first argue that there exists at least one resource between r and $t(p)$ in p. By assumption we have $\text{cost}_s(r) > \frac{1}{2}\text{cost}_s(r^*)$, where by definition r^* is such that $\text{cost}_s(r^*) \geq \text{cost}_s(u)$ for every resource $u \in R^-$ such that there is a path from u to r in the directed graph induced by Ψ. Thus, property 3 implies that each arc of the form (r, w) with positive flow in Φ has $w \notin R^+$. We can the conclude that $w \neq t(p) \in R^+$. This implies $p(t(p)) > p(r) + 1$. By (4.23) we then have $\Omega(p, p(r)) \geq \frac{1}{2}\lambda(r)$ for all $p \in P(r) \setminus P^0(r)$. We have

$$\frac{A_r}{B_r} = \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p\in P(r)} \Omega(p, p(r)) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \sum_{p\in P(r)} \Omega(p, p(r)) + \sum_{p\in P^0(r)} \text{cost}_s^+(t(p)) - \frac{a_r}{2}\text{cost}_s^+(r)}$$

$$\leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{p\in P^0(r)} \Omega(p, 0) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{p\in P^0(r)} (\Omega(p, 0) + \text{cost}_s^+(t(p))) - \frac{a_r}{2}\text{cost}_s^+(r)} \tag{4.32}$$

$$\leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{i=1}^{3}\sum_{p\in P_i^0(r)} \Omega(p, 0) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{i=1}^{3}\sum_{p\in P_i^0(r)} \left(\Omega(p, 0) + \frac{1}{2^{i-1}}\text{cost}_s(r)\right) - \frac{a_r}{2}\text{cost}_s^+(r)} \tag{4.33}$$

Inequality (4.32) holds because $\Omega(p, p(r)) \geq \frac{1}{2}\lambda(r)$ for all $p \in P(r) \setminus P^0(r)$ and $\frac{A_r}{B_r} \geq 1$. Inequality (4.33) holds because we have $P^0 = P_1^0 \cup P_2^0 \cup P_3^0$.

To derive an upper bound on the right-hand-side of (4.33), we do the following. First, for every $p \in P_1^0(r)$, we subtract $\Omega(p, 0) - \frac{3}{4}\text{cost}_s(r)$ from the numerator and subtract $\Omega(p, 0)$ from the denominator. Because $\Omega(p, 0) \geq 0$ for all $p \in P_1^0(r)$ by inequality (4.24) and $\frac{A_r}{B_r} \geq 1$, this yields an upper bound. Secondly, for every $p \in P_2^0(r)$, we subtract $\Omega(p, 0) - \frac{3}{4}\text{cost}_s(r)$ from the numerator and we subtract $\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r)$ from the denominator. This will also produce an upper bound, because $\Omega(p, 0) - \frac{1}{2}\text{cost}_s(r) \geq 0$ by inequality (4.25) and $\frac{A_r}{B_r} \geq 1$. Finally, for every $p \in P_3^0(r)$, we subtract $\Omega(p, 0) - \frac{3}{4}\text{cost}_s(r)$ from both the numerator and the denominator. This will yield an upper bound because $\Omega(p,0) - \frac{3}{4}\text{cost}_s(r) \geq 0$ by inequality (4.26) and $\frac{A_r}{B_r} \geq 1$. Then from (4.33) we have

$$\frac{A_r}{B_r} \leq \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{p\in P^0(r)} \frac{3}{4}\text{cost}_s(r) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{1}{2}\lambda(r) + \sum_{p\in P^0(r)} \text{cost}_s(r) - \frac{a_r}{2}\text{cost}_s^+(r)} \tag{4.34}$$

$$= \frac{\frac{(s_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{a_r}{2}(\frac{s_r}{s_r+1} - \frac{1}{2})\text{cost}_s^+(r) + \sum_{p\in P^0(r)} \frac{3}{4}\frac{(s_r)}{s_r+1}\text{cost}_s^+(r) - \frac{a_r}{2}\text{cost}_s^+(r)}{\frac{(q_r)^2}{s_r+1}\text{cost}_s^+(r) + \frac{a_r}{2}(\frac{s_r}{s_r+1} - \frac{1}{2})\text{cost}_s^+(r) + \sum_{p\in P^0(r)} \frac{(s_r)}{s_r+1}\text{cost}_s^+(r) - \frac{a_r}{2}\text{cost}_s^+(r)}, \tag{4.35}$$

where the equation (4.35) follows from (4.9) and $\lambda(r) = \text{cost}_s(r) - \frac{1}{2}\text{cost}_s^+(r)$. Finally, we rewrite the right-hand-side of (4.35) by factoring out $\text{cost}_s^+(r)$ and exploiting $|P^0(r)| = b_r$, $s_r = a_r + b_r + z_r$, and $q_r = a_r + z_r$. Next, we derive an upper bound by considering the maximum over all possible values of a_r, b_r and z_r. We obtain

$$\frac{A_r}{B_r} \leq \frac{\frac{(a_r+b_r+z_r)^2}{a_r+b_r+z_r+1} + \frac{a_r}{2}(\frac{a_r+b_r+z_r}{a_r+b_r+z_r+1} - \frac{1}{2}) + b_r\frac{3}{4}\frac{(a_r+b_r+z_r)}{a_r+b_r+z_r+1} - \frac{a_r}{2}}{\frac{(a_r+z_r)^2}{a_r+b_r+z_r+1} + \frac{a_r}{2}(\frac{a_r+b_r+z_r}{a_r+b_r+z_r+1} - \frac{1}{2}) + b_r\frac{(a_r+b_r+z_r)}{a_r+b_r+z_r+1} - \frac{a_r}{2}} \tag{4.36}$$

$$\leq \max_{a,z\geq 0, b\geq 1} \frac{\frac{(a+b+z)^2}{a+b+z+1} + \frac{a}{2}(\frac{a+b+z}{a+b+z+1} - \frac{1}{2}) + \frac{3}{4}b\frac{(a+b+z)}{a+b+z+1} - \frac{a}{2}}{\frac{(a+z)^2}{a+b+z+1} + \frac{a}{2}(\frac{a+b+z}{a+b+z+1} - \frac{1}{2}) + b\frac{(a+b+z)}{a+b+z+1} - \frac{a}{2}} \leq \frac{17}{7}. \tag{4.37}$$

This completes the proof of the theorem. $\square$

4.5 Conclusion

We have investigated the impact of matroid structures on the PoA of symmetric congestion games. In the symmetric case, the PoA of general matroid congestion games is still not completely understood. For graphic matroids and $N = 2, 3, 4$ or infinity with affine delay functions, the PoA can be as large as the worst-case PoA of symmetric congestion games, which is equal to $\frac{5N-2}{2N+1}$ [23]. However, for arbitrary N or different delay functions we don't know whether the the worst-case PoA of symmetric congestion games can be achieved by symmetric matroid congestion games. Our results indicate that if we restrict to paving matroid, the worst-case PoA is significantly smaller than that of symmetric congestion games. A similar result had been previously established by De Jong et al. [18] for k-uniform matroids and affine delays. However, k-uniform matroids are only a mild generalization of singleton congestion games. Paving matroid, on the other hand, are a substantial generalization of k-uniform matroids, since they are conjectured to represent the vast majority of matroids. Since paving matroid are quite more complex that k-uniform matroids, it is not as easy to characterize the worst-case PoA. There is still a gap between our upper and lower bounds, and we conjecture that the our upper bounds are not tight (see Remarks 4.11 and 4.14) .

Our approach to bound the PoA relies on a constant α that we have quantified for both k-uniform matroids and paving matroids (Theorem 4.7). In particular, we can set $\alpha = z(\mathcal{D})$ for k-uniform matroids and $\alpha = z(\mathcal{D})^2$ for paving matroids. Since paving matroids of rank k contain circuits whose size is smaller than the circuit size of k-uniform matroids, this suggests that the difference between the sizes of bases and circuits might impact the PoA. Let δ be a parameter that is equal to the rank of the matroid minus the size of the smallest circuit in the matroid. We conjecture that for $\delta \geq 0$ we can satisfy the assumptions of Theorem 4.7 with $\alpha = z(\mathcal{D})^{2(\delta+1)}$. Thus, we would get an upper bound on the PoA which is equal to $\rho(\mathcal{D})z(\mathcal{D})^{2(\delta+1)}$. For polynomial delays of highest degree p, this bound is in $O((C^p)(p/lnp))$, where $C = 4^{\delta+1}$. For fixed δ and large p this bound is still better than the PoA of general congestion games, that is in $O((p/lnp)^{p+1})$. To summarize, it is possible that our approach could be extended to upper bound the PoA in arbitrary matroid congestion games where we have an upper bound on δ. On the other hand, our approach might fail to provide meaningful upper bounds for small values of p or when the circuits can be much smaller than the rank. Besides the *size* of the circuits, we suspect that the way in which the circuits overlap can affect the PoA. For example, circuits of k-uniform matroids are highly symmetric. When dealing with paving matroids, we observed that instances with highly symmetric circuits displayed a lower PoA. On the other hand, the paving matroid congestion game example in Section 4.3, whose PoA is larger than the worst-case PoA of uniform matroid congestion games, has circuits that more often overlap on a single resource. In conclusion, it is open to find lower and upper bounds of symmetric matroid congestion games that depend on the size of the matroid circuits and/or on their degree of symmetry.

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

[2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., USA, 1993.

[3] 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.

[4] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, 2011.

[5] 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.

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

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

[8] 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), Oct. 2014.

[9] V. Bilò and C. Vinci. On the Impact of Singleton Strategies in Congestion Games. In K. Pruhs and C. 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.

[10] J. E. Blackburn, H. H. Crapo, and D. A. Higgs. A catalogue of combinatorial geometries. *Mathematics of Computation*, 27(121):155–s95, 1973.

[11] R. A. Brualdi. Comments on bases in dependence structures. *Bulletin of the Australian Mathematical Society*, 1(2):161–167, 1969.

[12] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part I*, ICALP'06, page 311–322, Berlin, Heidelberg, 2006. Springer-Verlag.

[13] 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.

[14] 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, Nov. 2019.

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

[16] A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. *SIAM Journal on Computing*, 39(5):1957–1987, 2010.

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

[18] J. de Jong, W. Kern, B. Steenhuisen, and M. Uetz. The asymptotic price of anarchy for k-uniform congestion games. In *Approximation and Online Algorithms*, pages 317–328, Cham, 2018. Springer International Publishing.

[19] J. de Jong, M. Klimm, and M. Uetz. Efficiency of equilibria in uniform matroid congestion games. In R. S. Martin Gairing, editor, *9th International Symposium on Algorithmic Game Theory (SAGT)*, volume 9928 of *Lecture Notes in Computer Science*, pages 105–116, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[20] A. Del Pia, M. Ferris, and C. Michini. Totally unimodular congestion games. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 577–588. SIAM, 2017.

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

[22] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of STOC '04*, 2004.

[23] W. Fokkema. The price of anarchy for matroid congestion games. ICALP 2023, workshop on congestion games, 2023.

[24] D. Fotakis. Stackelberg strategies for atomic congestion games. In L. Arge, M. Hoffmann, and E. Welzl, editors, *Algorithms – ESA 2007*, pages 299–310, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

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

[26] D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory of Computing Systems*, 47(1):113–136, 2010.

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

[28] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing nash equilibria for scheduling on restricted parallel links. *Theory of Computing Systems*, 47(2):405–432, 2010.

[29] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. *Journal of Computer and System Sciences*, 74(7):1199–1225, 2008.

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

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

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

[33] P. Kleer and G. Schäfer. Potential function minimizers of combinatorial congestion games: Efficiency and computation. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 223–240, New York, NY, USA, 2017. Association for Computing Machinery.

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

[35] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.

[36] Y. Kraakman. The price of anarchy of symmetric and semi-symmetric uniform congestion games, October 2021.

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

[38] 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.

[39] 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.

[40] M. Mavronicolas and P. Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007.

[41] D. Mayhew, M. Newman, D. Welsh, and G. Whittle. On the asymptotic proportion of connected matroids. *European Journal of Combinatorics*, 32(6):882–890, 2011. Matroids, Polynomials and Enumeration.

[42] D. Mayhew and G. F. Royle. Matroids with nine elements. *Journal of Combinatorial Theory, Series B*, 98(2):415–431, 2008.

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

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

[45] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

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

[47] J. Oxley. *Matroid Theory*. Oxford University Press, 02 2011.

[48] R. Pendavingh and J. van der Pol. On the number of matroids compared to the number of sparse paving matroids. *The Electronic Journal of Combinatorics*, 22(2), 2015.

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

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

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

[52] 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.

[53] 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.

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

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

[56] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

[57] A. Skopalik and B. Vöcking. Inapproximability of pure Nash equilibria. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 355–364, New York, NY, USA, 2008. ACM.

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

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