



Disjoint Paths, Dynamic Equilibria, and the Design of Networks

Marcus Raphael Kaiser

Vollständiger Abdruck der von der Fakultät für Mathematik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzende: Prof. Dr. Silke Rolles
Prüfende der Dissertation: 1. Prof. Dr. Andreas S. Schulz
2. Prof. José R. Correa, Ph.D.
(Universidad de Chile)

Die Dissertation wurde am 20.10.2021 bei der Technischen Universität München eingereicht und durch die Fakultät für Mathematik am 02.03.2022 angenommen.

Abstract

We study three network flow models that are governed by the interaction of the flow on links. Besides two models that exhibit negative effects of congestion, we investigate a third case that benefits from sharing links. The computational complexity of finding solutions to the resulting problems is addressed. Moreover, we look into the efficiency of solutions.

In the first problem that we examine, congestion needs to be avoided completely. The *disjoint shortest paths problem* asks for disjoint paths between given pairs of sources and sinks that are respectively shortest for given lengths. We deal with the case of two paths in an undirected network. A suitable decomposition allows us to devise a first polynomial-time algorithm for this setting that can handle zero-length edges in addition to positive lengths.

In *flows over time* under the *fluid queuing model*, the congestion of links is allowed and leads to the formation of queues. Waiting induces delay on flow traversing a link additional to a constant given transit time. Assuming selfish agents, we focus on the computational problem of finding *dynamic equilibria* in networks with a single source and a single sink. Our insights into the structure of the associated *thin flows with resetting* allow us to devise a first constructive algorithm for right-monotone inflow rate. Further, we use our findings to provide a polynomial-time algorithm for the computation of thin flows with resetting in series-parallel networks.

In contrast to both these settings, the collective usage of links is encouraged in the third discussed model. In a *network cost-sharing game*, a number of selfish agents selects paths between their respective sources and sinks. We deal with undirected networks. The players incur a cost for the usage of an edge that is nonincreasing in the number of players that are using it. We extend several results for constant edge costs to a large class of concave cost functions. On the one hand, we obtain extensive hardness results for the computation of *Nash equilibria*. On the other hand, we make progress in the determination of the *price of stability*, particularly, for the class of *broadcast games*.

Zusammenfassung

Diese Arbeit behandelt drei Netzwerkflussmodelle, welche sich durch die Interaktion von Fluss auf den Kanten auszeichnen. Neben zwei Modellen, in denen die Überlastung von Kanten zu negativen Effekten führt, wird ein drittes Szenario betrachtet, in dem es sich lohnt Kanten gemeinsam zu nutzen. Alle drei werden hinsichtlich der Komplexität untersucht, eine Lösung algorithmisch zu berechnen. Darüberhinaus wird auf den Wirkungsgrad der Lösungen eingegangen.

Die erste untersuchte Problemstellung schließt eine Überlastung von Kanten von Grund auf aus. Beim *Problem der disjunkten kürzesten Wege* gilt es, disjunkte Wege zwischen gegebenen Paaren von Quellen und Senken zu finden, welche bezüglich der gegebenen Kantenlängen jeweils kürzeste Wege sind. Im Mittelpunkt dieser Arbeit stehen zwei disjunkte Wege in einem ungerichteten Netzwerk. Eine geeignete Zerlegung ermöglicht deren Berechnung in polynomieller Zeit, wobei neben positiven Längen auch Kanten der Länge Null erlaubt sind.

Das *Fluid Queuing Model* lässt die Überlastung von Kanten durch *dynamischen Fluss* zu. Überschüssiger Fluss sammelt sich in Warteschlangen. Das Warten in diesen führt zur Verzögerung des Flusses beim Überqueren von Kanten, welche die gegebenen konstanten Durchflusszeiten erhöhen. Unter der Annahme, dass jeder Flusspartikel egoistisch handelt, wird im Speziellen die algorithmische Berechnung eines *dynamischen Equilibriums* in einem Netzwerk mit nur einer Quelle und einer Senke betrachtet. Erkenntnisse zur Struktur der zu Grunde liegenden *Thin Flows with Resetting* erzielen einen ersten Algorithmus für eine rechtsmonotone Einflussrate. Gleichzeitig führen sie zu einem polynomiellen Algorithmus zur Berechnung von Thin Flows with Resetting in serienparallelen Netzwerken.

Im Gegensatz zu den beiden ersten, bietet im dritten Modell die gemeinsame Nutzung von Kanten einen Vorteil. In einem *Network Cost-Sharing Game* wählt jeder Spieler einen Weg, um seine Quelle mit seiner Senke zu verbinden. Für die Verwendung einer Kante fallen dabei Kosten für jeden Spieler an, welche monoton in der Anzahl der sie verwendenden Spieler fällt. Zahlreiche Ergebnisse für konstante Kantenkosten werden auf eine ganze Klasse von konkaven Kostenfunktionen erweitert. Einerseits werden weitreichende Resultate zur Schwere der Berechnung von *Nash Equilibrien* vorgestellt. Andererseits wird die Bestimmung des *Price of Stability* vor allem für die Klasse der *Broadcast Games* vorangetrieben.

Contents

1	Introduction	1
2	Preliminaries	5
2.1	Basic Notation and Terminology	5
2.2	Graphs	7
2.3	Network Flows	11
2.4	Game Theory	12
2.5	Computational Problems and Complexity Theory	15
2.6	Linear Programming	18
2.7	Linear Complementarity Problems	21
3	Two Disjoint Shortest Paths	23
3.1	Introduction	24
3.1.1	Related Literature	24
3.1.2	Our Contribution	27
3.2	A Decomposition of Disjoint Paths	27
3.3	Disjoint Paths in Weakly Acyclic Mixed Graphs	30
3.4	Undirected Disjoint Shortest Paths	32
3.4.1	Orienting Shortest Paths	32
3.4.2	Disjoint Paths in the Partial Orientation	35
3.5	Closing Remarks	38
4	Dynamic Equilibria under the Fluid Queuing Network	39
4.1	Introduction	40
4.1.1	Related Literature	40
4.1.2	Our Contribution	43
4.2	The Fluid Queuing Model and Dynamic Equilibria	43
4.3	Normalized Thin Flows with Resetting	48
4.3.1	A Linear Complementarity Problem	49
4.3.2	Parametric Normalized Thin Flows with Resetting	57

Contents

4.4	Evolution of Dynamic Equilibria	60
4.4.1	A Differential Equation	60
4.4.2	Dynamic Equilibria for Right-Monotone Inflow	64
4.5	Thin Flows with Resetting in Series-Parallel Graphs	70
4.6	Closing Remarks	74
5	Nash Equilibria in Network Cost-Sharing Games	75
5.1	Introduction	76
5.1.1	Related Literature	79
5.1.2	Our Contribution	83
5.2	Structure of Nash equilibria	84
5.3	Computational Complexity of Nash Equilibria	89
5.3.1	Formulations	89
5.3.2	Intractability Results	94
5.3.2.1	PLS-Hardness of Computing Nash Equilibria	95
5.3.2.2	NP-Hardness of Computing a Minimum-Cost Nash Equilibrium	100
5.3.2.3	NP-Hardness of Computing a Global Potential Minimizer	101
5.3.2.4	Slowly Improving Dynamics	103
5.4	Efficiency of Nash Equilibria	105
5.4.1	The Price of Anarchy	105
5.4.2	Upper Bounds on the Price of Stability	107
5.4.2.1	The Potential Function Method	107
5.4.2.2	The Homogenization-Absorption Framework	112
5.4.3	Lower Bounds on the Price of Stability	131
5.4.3.1	The Fan Graph	132
5.4.3.2	A Lower Bound for Constant Total Edge Cost	141
5.4.3.3	A Lower Bound for Affine Total Edge Cost	146
5.4.3.4	A Lower Bound for Polynomial Total Edge Cost	150
5.5	Closing Remarks	151
	Notation	155
	Index	159
	Bibliography	163

Chapter 1

Introduction

Network flows are ubiquitous. Be it traffic that is routed through a road or rail network, data which is transmitted via a communication network, or current that runs through an electrical network. All these applications have a common abstraction: one or more commodities **flow** through a network of nodes interconnected by links. Different assumptions on the network and the commodities are required in different applications. Cumulatively, they try to capture the fundamental behavior of the examined system. In this thesis, we examine three models which are determined to a large extent by the interaction between flows in a single network. The common usage of a link by multiple commodities leads to congestion which can result in various effects. A typical negative effect that comes to mind is **delay**, for example, caused by traffic jams. In this scenario, it is evident that a suitable widespread distribution of the flow over the whole network has the potential of reducing the congestion and, therefore, benefiting the whole system. In other settings, the negative effect might be even more severe to the point where congestion results in disaster. In an application, such a devastating interaction is represented, for example, by collision in the routing of automated machines. In this case, the primary goal is to send flow along **disjoint** routes. On the other hand, the collective usage of a link does not always imply impairment. Positive effects arise, for example, in the context of network design. If the construction of links exhibits economies of scale, multiple parties benefit from **sharing** the cost of commonly created infrastructure. Here, it is advantageous to route flow on overlapping routes.

The balance between a model's level of detail and its resulting complexity is necessary to keep it meaningful and its analysis tractable at the same time. We discuss basic assumptions and typical simplifications which are relevant in the scope of this thesis. The maybe most basic property of a network is its directionality. An **undirected** link represents a symmetric relation between its endpoints. Flow can traverse it in both directions. **Directed** links on the other hand can only be traversed in a single direction. An example are one-way roads. Typically, directed networks exhibit a higher complexity than undirected ones. Finding a path through a city center with narrow one-way streets is often more challenging than routing without having to worry about the allowed directions.

Another fundamental characteristic of networks is the capacity of its links. Physical networks often possess a natural limit to the amount of flow that a single link can carry. Virtual commodities, like currencies, on the other hand are not subject to such a restriction. If the links are **uncapacitated**, every flow particle with the same origin and

destination can take the same route through the network. By imposing **capacities** on a flow, a potentially much larger number of different routes is necessary.

While most physical commodities have an **atomic** size they can be split into, indivisibility is more crucial in some models than in others. Requiring the quantities to be **integral** generally can increase the intricacy of mathematical models immensely. Hence, the simplifying assumption of being **nonatomic**, that is, arbitrarily divisible is made if possible. This is more easily accepted when it comes to divisible goods like fluids, data, or currencies. For indivisible goods, for instance cars, this assumption can still be justified on a macroscopic level. Then an individual unit of flow only contributes a negligible amount. On a microscopic level with only a small amount of atomic units of flow, however, the requirement of integral quantities often is indispensable.

The complexity of network flows also highly depends on the number of commodities that share a single network. Here, **multiple commodities** are not necessarily of a different sort, but typically differ in the location of their sources and sinks in the network. The number of commodities in the model not necessarily has to match the number in its application. Multiple commodities that share a common source or sink, for example, can typically captured by a **single commodity** of flow in the model.

Another major feature of flow models, we want to highlight, is whether they include the dimension of time or not. In **static flow** models, capacity for a flow particle is reserved simultaneously on every link on its route. This model is suitable if timing does not play a role or the supply and demand structure stays constant over a long period of time. If this is not the case, however, **dynamic flows** or **flows over time** likely give a more appropriate model. Therein, flow spreads at a finite speed through the network. Further, they allow to capture temporal fluctuations in the supplies and demands.

For flow over time, additional choices in the model can be made. A central one is the possibility of flow to pause along its route. Depending on the application, there might be capacity for flow to collect at junctions or on links. In traffic models for instance, one could assume that flow clears junctions immediately. On links, however, there might be space for **queues**.

Maybe the most critical part of a model in terms of capturing the behavior and potential of a system is the objective. The most categorical question one can ask is whether a certain situation or structure can occur. This could be the existence of a flow with a given throughput. Such solutions demonstrate **feasibility** but often lack practicality. In an endeavor for efficiency, one is often interested in a solution that is particularly good. Quantifying a measure of quality by a single objective function allows to search for a best solution. This goal is pursued by the field of **optimization**. The implementation of such an optimal solution in practice, however, is not always realizable. In particular, it requires a central authority that can enforce all necessary decisions. Some scenarios, as for example individual transport, do not allow such an intervention. Here, the angle that **game theory** takes is more applicable. Under the assumption that the involved parties act selfishly, it asks for stable configurations called equilibria. Those are examined in terms of efficiency and the impact of centrally controllable parameters, like the charging of tolls.

This thesis covers three computational problems from different domains. Each highlights different properties of network flows and requires different solution techniques.

The **disjoint shortest paths problem** deals with completely avoiding congestion. It asks for the existence of disjoint paths in a network between given pairs of nodes which are respective shortest paths. Here, the length of a path is the sum of the given lengths of its links. The quality of a solution is ensured by the requirement of the paths being as short as they could be if no other commodities were present. We specifically treat two paths in an undirected network. In terms of network flows, we obtain two indivisible commodities. Every link provides the capacity for a single unit of flow only. We devise a first polynomial-time algorithm that is capable of deciding on the existence of two disjoint shortest paths under the presence of zero-length edges.

In flows over time under the **fluid queuing model**, a nonatomic flow of selfish particles travels through a directed network. The transit of a link takes flow a given constant amount of time. If the given capacity of a link is exceeded, a queue is formed at its entrance. Waiting in a queue adds additional delay for flow to traverse a link. Hence, congestion affects the system disadvantageously. We examine the case of flow traveling from a common source to a common sink, which makes a single commodity sufficient to describe all particles. We focus on the computational problem of equilibria in this model. Structural insights for a special type of static flows allows us to handle a large class of inflow rates. This provides the theoretic foundation for time-dependent supplies in this model.

Finally, we examine a model that exhibits positive impact of the common usage of links. In a **network cost-sharing game** a number of players connects their respective source and sink by a path in a network. The cost a player incurs for the usage of a link is nonincreasing in the number of users. We analyze these games in undirected networks. In the language of flows, every player constitutes a commodity a single indivisible unit of which needs to be routed through the network. We extend hardness results on the computation of equilibria. Moreover, we compare the quality of equilibria with optimal solutions which could be achieved by a centralized routing decision.

Chapter 2

Preliminaries

This chapter provides an overview of the notation and the terminology that is used throughout this thesis. Further, fundamental results that our findings build upon are stated. A listing of notation and an index of terminology can be found in the back matter.

2.1 Basic Notation and Terminology

Numbers. We denote the set of **natural numbers** (including zero) by \mathbb{N} . For the first $n \in \mathbb{N}$ positive natural numbers, we write $[n] := \{1, \dots, n\}$. To the set of **real numbers**, we refer as \mathbb{R} . The negative, nonpositive, nonnegative, and positive reals are denoted by $\mathbb{R}_{<0}, \mathbb{R}_{\leq 0}, \mathbb{R}_{\geq 0}, \mathbb{R}_{>0}$. The **positive** and **negative parts** of a real number $x \in \mathbb{R}$ are defined as $[x]_+ := \max\{x, 0\}$ and $[x]_- := \max\{-x, 0\}$, respectively. The n -th **harmonic number** is denoted by $H(n) := \sum_{k \in [n]} 1/k$.

Sets. We write $|X|$ for the **size** or **cardinality** of a finite set X . The family of all subsets of X with a given size $k \in \mathbb{N}$ is given by $\binom{X}{k} := \{Y \subseteq X \mid |Y| = k\}$. The **power set** of a finite set X is given by $2^X := \{Y \mid Y \subseteq X\}$. To emphasize that the union of two sets X and Y is disjoint, we write $X \uplus Y$. The **symmetric difference** of two sets X and Y is defined as $X \triangle Y := (X \setminus Y) \cup (Y \setminus X)$.

Relations. A **binary relation** R on a ground set U is a subset of $U \times U$. For $(u, v) \in R$, we write $u R v$. The **composition** $S \circ R$ of two binary relations S, R on U is defined as

$$\{(u, w) \in U \times U \mid \exists v \in U: u R v \wedge v S w\}.$$

Note that \circ is an associative operator; that is, $(Q \circ R) \circ S = Q \circ (R \circ S)$ holds for all binary relations Q, R , and S on U .

Vectors and matrices. We use $\mathbb{1}$ for the all-one vector of appropriate dimension. For a set of indices I , the vector $\mathbb{1}_I$ denotes the **characteristic vector** of the set I of appropriate dimension. For an index i , we write the **unit vector** of component i as $\mathbb{1}_i$. We use Id for the **identity matrix** of appropriate dimension. The **transposed** of a matrix $M \in \mathbb{R}^{m \times n}$ is denoted as $M^\top \in \mathbb{R}^{n \times m}$. The **submatrix** determined by the rows $I \subseteq [m]$ and the columns $J \subseteq [n]$ is written as $M_{I,J}$. If $I = [m]$ or $J = [n]$, we write

$M_{\bullet,J}$ and $M_{I,\bullet}$, respectively. Similarly, the **subvector** of a vector $x \in \mathbb{R}^n$ restricted to the components $I \subseteq [n]$ is denoted by x_I . We write the **uniform norm** of a vector x as $\|x\|_\infty$. Also, we use the short hand $x(I) := \sum_{i \in I} x_i$ for the sum of the components I of a vector x .

Functions. A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is **locally Lebesgue-integrable** if its Lebesgue-integral over any compact set of \mathbb{R} is finite. We denote the set of all locally Lebesgue-integrable functions on \mathbb{R} by $L^1_{\text{loc}}(\mathbb{R})$. As usual we turn to the equivalence classes of functions that are equal almost everywhere (a.e.). The **essential infimum** and **essential supremum** of $f \in L^1_{\text{loc}}(\mathbb{R})$ are defined as

$$\text{ess inf}(f) := \sup\{x \in \mathbb{R} \mid f \geq x \text{ a.e.}\} \quad \text{and} \quad \text{ess sup}(f) := \inf\{x \in \mathbb{R} \mid f \leq x \text{ a.e.}\}.$$

Mediant. Let $k \in \mathbb{N}$ numerators and denominators be given by $x_i \in \mathbb{R}_{\geq 0}, i \in [k]$ and $y_i \in \mathbb{R}_{> 0}, i \in [k]$, respectively. The **mediant** of the fractions $x_i/y_i, i \in [k]$ is defined by

$$\frac{\sum_{i \in [k]} x_i}{\sum_{i \in [k]} y_i}.$$

It fulfills the **mediant inequality**

$$\min_{i \in [k]} \frac{x_i}{y_i} \leq \frac{\sum_{i \in [k]} x_i}{\sum_{i \in [k]} y_i} \leq \max_{i \in [k]} \frac{x_i}{y_i}.$$

The bounds hold with equality if and only if the values of the fractions x_i/y_i are all the same.

Weighted harmonic mean. The **weighted harmonic mean** of $k \in \mathbb{N}$ nonnegative real numbers $x_i \in \mathbb{R}_{\geq 0}, i \in [k]$ with respect to nonnegative weights $w_i \in \mathbb{R}_{\geq 0}$ is defined by

$$\frac{\sum_{i \in [k]} w_i}{\sum_{i \in [k]} \frac{w_i}{x_i}}.$$

If there is an x_i which is zero, the weighted harmonic mean is set to zero. In the case, that the total weight $\sum_{i \in [k]} w_i$ is zero, it is not defined. Similarly to the mediant, the weighted harmonic mean satisfies

$$\min_{i \in [k]} x_i \leq \frac{\sum_{i \in [k]} w_i}{\sum_{i \in [k]} \frac{w_i}{x_i}} \leq \max_{i \in [k]} x_i.$$

Further, it has the following monotonicities. The weighted harmonic mean of x_1 and x_2 with respect to nonnegative weights w_1, w_2 is nondecreasing in x_1 and x_2 (strictly if the respective weight is positive). If $x_1 \leq x_2$, it is nonincreasing in w_1 and nondecreasing in w_2 (strictly if $x_1 < x_2$ and respectively $w_2 > 0$ or $w_1 > 0$). Note that these monotonicities

extend to more than two elements, because the harmonic mean is associative in the sense that, for $\emptyset \neq I \subsetneq [k]$, we have

$$\frac{\sum_{i \in [k]} w_i}{\sum_{i \in [k]} \frac{w_i}{x_i}} = \frac{\sum_{i \in I} w_i + \sum_{i \in [k] \setminus I} w_i}{\sum_{i \in I} w_i \left(\frac{\sum_{i \in I} w_i}{\sum_{i \in I} \frac{w_i}{x_i}} \right)^{-1} + \sum_{i \in [k] \setminus I} w_i \left(\frac{\sum_{i \in [k] \setminus I} w_i}{\sum_{i \in [k] \setminus I} \frac{w_i}{x_i}} \right)^{-1}}.$$

This means that the weighted harmonic mean of k elements can be viewed as the weighted harmonic mean of the weighted harmonic means in any partition.

2.2 Graphs

Many problems in operations research and related fields involve some kind of network, be it traffic networks, broadcast networks, social networks, or others. Networks are usually modeled using **graphs**. Entities in the network, like routing points or junctions, are represented by **vertices**. Pairwise relations between them, like connections or links, are modeled via **arcs** or **edges**. Arcs are used if the relation exhibits directionality, like links which can only be used in one way. Edges on the other hand represent symmetric relations. Due to this fundamental difference of arcs and edges, graphs are usually categorized by their directionality into **undirected**, **directed**, and **mixed** graphs.

Undirected graphs. An **undirected graph** is given by a tuple $G = (V, E)$ where V is a finite set of vertices and $E \subseteq \binom{V}{2}$ is a set of edges. We also write $V(G)$ and $E(G)$ for the vertices and edges of the graph G , respectively (if these sets are not explicitly named). The set of vertices which share an edge with a vertex $v \in V$ are called its **neighbors** and denoted by

$$N_E(v) := \{w \in V \mid \{v, w\} \in E\}.$$

For a subset of vertices $\emptyset \neq U \subsetneq V$, the set

$$\delta_E(U) := \{e \in E \mid |e \cap U| = 1\}$$

denotes the set of edges in E that have exactly one of their vertices in U . In this context, the term **cut** is used to refer to U and sometimes also to $\delta_E(U)$. If U contains only a single vertex u , we write $\delta_E(u)$ instead of $\delta_E(\{u\})$ to denote the set of edges in E that are **incident** to u .

For two vertices $v, w \in V$, a v - w **path** in G is a subset $P \subseteq E$ of edges along a sequence of vertices $v = v_0, v_1, \dots, v_l = w$, that is,

$$P = \{\{v_{i-1}, v_i\} \in E \mid i \in [l]\}.$$

A path is called **simple** if the vertices v_0, \dots, v_l are pairwise different. In the case $v = w$, a v - w path is called **cycle**. A cycle is still called **simple** if the vertices $v_1, \dots, v_l = v_0$ are pairwise different. If there exists a v - w path, we also say that v and w are **connected**. The **connected components** of G are given by the partition of its vertices into subsets

such that two vertices are in the same subset if and only if they are connected in G . If there is only one connected component, the graph G is called **connected**.

Directed graphs. A **directed graph** is a tuple $G = (V, A)$ where V again is a finite set of vertices and $A \subseteq V \times V$ is a set of arcs. The sets of vertices and arcs of G are also referred to as $V(G)$ and $A(G)$, respectively. We distinguish the related vertices of a vertex v by the direction of the corresponding arcs. The **neighbors** $N_A(v)$ of v are partitioned into the sets of **in-neighbors** and **out-neighbors** which are denoted respectively by

$$N_A^-(v) := \{u \in V \mid (u, v) \in A\} \quad \text{and} \quad N_A^+(v) := \{w \in V \mid (v, w) \in A\}.$$

Similarly, the set of arcs in a cut $\emptyset \neq U \subsetneq V$ is partitioned into the sets of **incoming** and **outgoing** arcs, respectively denoted by

$$\delta_A^-(U) := \{(v, w) \in A \mid v \notin U, w \in U\} \quad \text{and} \quad \delta_A^+(U) := \{(v, w) \in A \mid v \in U, w \notin U\}.$$

Their union is denoted by $\delta_A(U)$. Again, we use the shorthand $\delta_A^-(u)$, $\delta_A^+(u)$, and $\delta_A(u)$ if U consists of a single vertex u .

For two vertices $v, w \in V$, a **(directed) v - w path** in G is a subset $P \subseteq A$ of arcs that follows a sequence of vertices $v = v_0, v_1, \dots, v_l = w$, that is,

$$P = \{(v_{i-1}, v_i) \in A \mid i \in [l]\}.$$

If there exists a directed v - w path, we also say that w can be **reached** from v . The notions of **(directed) cycles** as well as simple paths and cycles are defined analogously to the undirected case. The **weakly connected components** of G are the connected components of the undirected graph resulting from dropping the direction of all arcs. In contrast to that, the **strongly connected components** of G are the partition of its vertices into subsets such that two vertices are in the same subset if and only if they can reach each other. If a graph G has only one weakly or strongly connected component it is called **weakly** or **strongly connected**, respectively.

Mixed graphs. A **mixed graph** is a tuple $G = (V, \mathcal{A})$ where V is the set of vertices and $\mathcal{A} = A \cup E$ is the union of arcs A and edges E . For technical reasons, we assume that never both $\{v, w\} \in E$ and $(v, w) \in A$ are present for any $v, w \in V$. We use the notations for undirected and directed graphs in order to express the neighborhoods and the incidence with respect to E and A . All neighbors of a vertex $v \in V$ are denoted by $N_{\mathcal{A}}(v) := N_E(v) \cup N_A(v)$. Similarly, the set of its incident arcs and edges is denoted by $\delta_{\mathcal{A}}(v) := \delta_E(v) \cup \delta_A(v)$. A v - w **path** $P \subseteq \mathcal{A}$ in G for vertices $v, w \in V$ is determined by vertices v_0, \dots, v_l such that either $\{v_{i-1}, v_i\} \in E$ or $(v_{i-1}, v_i) \in A$ holds for all $i \in [l]$. P then is the set of all such arcs and edges. The concepts of simplicity and cycles generalize naturally. Also weakly and strongly connected components can be defined for mixed graphs based on this definition of paths.

Relations and operations. The following notions apply to graphs $G = (V, \mathcal{A})$ of any type. Two paths are **vertex-disjoint** or **arc/edge-disjoint** if their sets of vertices, or their sets of arcs and edges are disjoint, respectively. A v - w path is **internally vertex-disjoint** to a set of vertices $U \subseteq V$ if its vertex set intersects U only in v or w . The subgraph of G **induced** by a subset of vertices $U \subseteq V$ is the graph $G[U]$ that results from G by deleting all vertices $V \setminus U$ and the incident arcs and edges thereof. The **contraction** G/U of a set of vertices U , in some sense, is a complementary operation to that. We define G/U as the graph that results from G by joining all vertices in U into a single new vertex, which inherits all neighbors of vertices in U . Arcs and edges within U are removed by the contraction. For pairwise disjoint vertex sets $U_1, \dots, U_k \subseteq V$, we denote by $G/\{U_1, \dots, U_k\}$ the graph that results from G by contracting U_1, \dots, U_k into k vertices. Note that the resulting graph is independent of the order of the contraction of U_1, \dots, U_k .

Algebraic representation. We use the following two algebraic representations of a directed graph $G = (V, A)$. Similar concepts exist for undirected graphs as well. The **incidence matrix** $B \in \{-1, 0, 1\}^{V \times A}$ of G is defined by

$$B_{v,a} = \begin{cases} -1 & \text{if } a \in \delta_A^+(v) \\ 1 & \text{if } a \in \delta_A^-(v) \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } v \in V, a \in A.$$

Its positive and negative parts are denoted by B^+ and B^- , respectively. In particular, $B = B^+ - B^-$. Note that the graph G is completely determined by its incidence matrix.

For arc weights $\nu \in \mathbb{R}^A$, the **weighted Laplacian matrix** $L \in \mathbb{R}^{V \times V}$ of G is defined by its entries

$$L_{v,w} = \begin{cases} \nu(\delta^-(v)) & \text{if } v = w \\ -\nu_{v,w} & \text{if } (v, w) \in A \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } v, w \in V.$$

The graph G and the weights ν can be fully recovered from the weighted Laplacian L . For our purposes, the weights ν are chosen to be the capacities of the arcs.

The following relation between the incidence matrix and the weighted Laplacian matrix is straightforward to check. It holds

$$L = BD(B^+)^T$$

where $D = \text{diag}(\nu)$ is the square matrix with diagonal ν and zero off-diagonal entries.

Trees. Let $G = (V, E)$ be an undirected graph. A set $T \subseteq E$ is called a **tree** if it is connected and does not contain any cycles. A tree T is **spanning** G if $V(T) = V$. For $v, w \in V(T)$, we denote the unique v - w path in T by $T[v, w]$.

A **rooted tree** is a tuple (T, r) such that T is a tree and $r \in V(T)$. If r is clear from the context, we identify T with (T, r) . The **depth** of a vertex $v \in V$ in T then is $|T[v, r]|$. The **height** of T is the maximal depth of any vertex. With a vertex $v \in V \setminus \{r\}$, in a rooted tree (T, r) we associate the edge $e_T(v) \in \delta(v) \cap T[r, v]$ being the first edge on the path from v to r . Note that $e_T(v)$ is uniquely determined by T , r , and v .

The set of **ancestors** $A_T(v)$ of a vertex v in T is the set of vertices appearing on the path $T[r, v]$ (including v). Similarly, the set of **descendants** $D_T(v)$ of a vertex v in T is the set of vertices u that v is an ancestor of, that is,

$$D_T(v) = \{u \in V \mid v \in V(T[u, r])\}.$$

For $v, w \in V$, the **lowest common ancestor** $\text{lca}_T(v, w)$ in T is defined as the unique vertex which is an ancestor of v and w , and has the largest depth $|T[r, \text{lca}_T(v, w)]|$. For a vertex v , the **subtree of T rooted at v** is the induced tree $T[D_T(v)]$.

We extend the notions of ancestors and descendants to sets of vertices $W \subseteq V$. Denote by $D_T(W) := \bigcup_{w \in W} D_T(w)$ the union of the descendants of any vertex in W . For such a descendant $u \in D_T(W)$, the **lowest ancestor** $a_T^W(u)$ of u in W is the vertex of maximal depth in $W \cap A_T(u)$. In order to express the inverse relation of a_T^W , define for $w \in W$ the set of **direct descendants** of w with respect to W as

$$D_T^W(w) := \{u \in D_T(w) \mid w = a_T^W(u)\}.$$

Note that $\bigcup_{w \in W} D_T^W(w)$ is a partition of the descendants $D_T(W)$ with respect to their lowest ancestor.

There is a corresponding concept to a rooted tree in directed graphs. For a root vertex r and a subset of vertices $U \subseteq V$, an **arborescence** is a set of arcs that contains exactly one directed r - v path for every $v \in U$. Thus, it is a tree with every edge oriented away from the root.

Graph classes. With trees we already have introduced a first class of graphs. We highlight three further classes which are relevant in the context of this thesis. For more details and further graph classes, we refer the reader to the textbook of Brandstädt, Le, and Spinrad (1999).

The class of **directed acyclic graphs**—as the name suggests—is given by directed graphs which do not contain any cycle. They induce natural orderings on their vertices. A total ordering of the vertices is called a **topological ordering** with respect to a directed acyclic graph if v precedes w in the ordering for every arc (v, w) .

Another well-known class of graphs are **planar graphs**. Those are graphs which can be drawn in the plane by representing its vertices and arcs or edges by points and arrows or lines, respectively, such that no pair of arrows/lines overlap. They appear naturally in applications that are limited to surfaces and do not allow the crossing of links.

The third relevant class for our purposes is given by **series-parallel directed graphs**. A directed graph is **two-terminal** if it has a unique source s and a sink vertex t ; that is, s and t are the unique vertices without any incoming and outgoing arcs, respectively.

The **series composition** of two two-terminal graphs G_1 and G_2 is obtained by taking their disjoint union and contracting the two sources and also the two sinks. We denote it by $G_1 * G_2$. The **parallel composition** of two two-terminal graphs G_1 and G_2 results from contracting the sink of G_1 and the source of G_2 in the disjoint union of the graphs. We write $G_1 \parallel G_2$ for it. The class of series-parallel directed graphs then is the smallest set of graphs which contains the graph(s) with two vertices and a single arc in between, and which is closed under series and parallel composition. From this recursive definition, it follows inductively that series-parallel directed graphs are acyclic and planar.

2.3 Network Flows

Typically networks are used for the conveyance of some sort of commodity. This can be for example vehicles in traffic networks or data in the case of communication networks. A **flow** describes the routing of a commodity through a network represented by a directed graph $G = (V, A)$. Formally, a flow is defined as a nonnegative vector $x \in \mathbb{R}_{\geq 0}^A$. The **excess** of x at a vertex $v \in V$ is given by the difference of incoming and outgoing flow

$$x(\delta_A^-(v)) - x(\delta_A^+(v)).$$

The flow x fulfills **weak flow conservation** at a vertex $v \in V$ if the excess at v is nonnegative. If it is zero, we say that x fulfills **strict flow conservation**. Typically, the arcs in a network have assigned capacities $\nu \in \mathbb{R}_{\geq 0}^A$ that limit the amount of flow they can carry. Then the flow x satisfies the **capacity constraints** if $x_a \leq \nu_a$ holds for all arcs $a \in A$.

In the most basic case of **singlecommodity flow**, we consider flow that emerges from a source $s \in V$ and needs to be routed to a sink $t \in V$. We generally assume that $s \neq t$. An s - t **flow** is a flow x that fulfills the capacity constraints and strict flow conservation at every nonterminal vertex $V \setminus \{s, t\}$. Its **value** $|x|$ is defined as the excess at the sink. In a more general setting, a vector $d \in \mathbb{R}^V$ defines the demands ($x_v > 0$) and supplies ($x_v < 0$) of a single commodity at the vertices in G . Then, the excess at a vertex v is defined as

$$x(\delta_A^-(v)) - x(\delta_A^+(v)) - d_v,$$

and flow conservation is imposed on all vertices.

Path flows. Flows can also be defined in a more global manner. Let \mathcal{P}_{st} be the set of all simple s - t paths in G and \mathcal{C} be the set of all simple cycles in G . Then an s - t **path flow** is a nonnegative vector $x \in \mathbb{R}_{\geq 0}^{\mathcal{P}_{st} \cup \mathcal{C}}$. Every such path flow x induces an (arc) flow in the sense of the preceding definition via the transformation

$$x_a = \sum_{\substack{P \in \mathcal{P}_{st}: \\ a \in P}} x_P + \sum_{\substack{C \in \mathcal{C}: \\ a \in C}} x_C \quad \text{for all } a \in A.$$

It can be seen that the induced arc flow $(x_a)_{a \in A}$ satisfies strict flow conservation at every nonterminal vertex $V \setminus \{s, t\}$. Its value is equal to $\sum_{P \in \mathcal{P}_{st}} x_P$. There is also a mapping in the opposite direction. Every s - t flow defined on the arcs of G satisfying strict flow conservation can be **decomposed** into a path flow. This decomposition, however, is not unique.

Multicommodity flows. The preceding describes the routing of a single commodity through a network. This can be generalized by simultaneously routing k commodities with individual source s_i and sink t_i for every $i \in [k]$ through a common network. A **multicommodity flow** is given by an s_i - t_i flow $x^{(i)}$ for every $i \in [k]$. The capacity of each arc in the network is shared by all commodities. Thus, a multicommodity flow fulfills the capacity constraints if $\sum_{i \in [k]} x_a^{(i)} \leq \nu_a$ holds for all arcs $a \in A$.

Flows over time. The definition of flow does not account for temporal aspects of conveyance. To emphasize this, we also speak of **static flow**. In contrast to that, a **flow over time** describes a commodity which travels through a network at a finite speed. The time it takes flow to traverse the arcs is specified by **transit times** $\tau \in \mathbb{R}_{\geq 0}^A$. Formally, a flow over time is defined as a function $f: \mathbb{R} \rightarrow \mathbb{R}^A$ which maps a point in time $\vartheta \in \mathbb{R}$ to a flow value $f_a(\vartheta)$ on every arc $a \in A$. Every component f_a is required to be locally Lebesgue-integrable. Usually, we assume that the network is initially empty and f vanishes on $\mathbb{R}_{<0}$. The **excess** of the flow over time f at time $\vartheta \in \mathbb{R}$ in a vertex $v \in V$ is given by

$$\sum_{a \in \delta_A^-(v)} \int_0^{\vartheta} f_a(\theta - \tau_a) d\theta - \sum_{a \in \delta_A^+(v)} \int_0^{\vartheta} f_a(\theta) d\theta.$$

f fulfills **weak** or **strict flow conservation** at a time ϑ in a vertex v if the excess at ϑ and v is nonnegative or zero, respectively. In a capacitated network, we say that f fulfills the **capacity constraints** if $f(\vartheta) \leq \nu$ holds for all $\vartheta \in \mathbb{R}$ (component-wise). The **time horizon** of a flow over time f is determined by the last point in time of nonzero flow, that is, the essential supremum of its support $\text{ess sup}\{\vartheta \in \mathbb{R} \mid f(\vartheta) \neq 0\}$. For a source $s \in V$ and sink $t \in V$, an s - t **flow over time** is a flow over time f which fulfills the capacity constraints and obeys strict flow conservation at every time $\vartheta \in \mathbb{R}$ and every nonterminal vertex $v \in V \setminus \{s, t\}$.

For more details on network flows, see the textbook of Ahuja, Magnanti, and Orlin (1993). An introductory article on flows over time is available by Skutella (2009).

2.4 Game Theory

In many real-world systems, there is no central authority that controls all decisions. In game theory, systems with multiple competing decision makers are examined. Each present decision maker is called **player** or **agent**. Players are generally assumed to act selfishly; that is, they want to optimize their individual objective function. Depending

on the application, this is phrased as the maximization of utility or the minimization of cost. Each player is restricted to play one of their available strategies. Their cost depends on the strategies chosen by all players.

Formally, such a system can be modeled as a **normal form game**. There are $n \in \mathbb{N}$ players. For each player $i \in [n]$, there is a finite set Σ_i of available **strategies**. A tuple consisting of a strategy of each player is called **strategy profile**. We denote the set of strategy profiles by $\Sigma = \times_{i \in [n]} \Sigma_i$. Every player $i \in [n]$ has a player cost function $C_i: \Sigma \rightarrow \mathbb{R}$. The goal of every player is to minimize her cost $C_i(\sigma)$ in a strategy profile $\sigma \in \Sigma$. However, she can only influence her choice of σ_i and not the strategies of the other players. A **unilateral deviation** of a player $i \in [n]$ in a strategy profile σ to a strategy $\sigma'_i \in \Sigma_i$ is denoted by (σ_{-i}, σ'_i) . It is called **improving** if $C_i(\sigma_{-i}, \sigma'_i) < C_i(\sigma)$. The strategy σ'_i is a **best response** of player i to σ_{-i} if it minimizes the cost $C_i(\sigma_{-i}, \sigma'_i)$. A strategy is called **dominant** if it is a best response to all strategy profiles. The prevailing solution concept for such games in the literature is introduced by Nash (1950). A strategy profile is considered stable if every player's strategy is a best response. Then no player can improve her cost by unilaterally deviating. More formally, a strategy profile $\sigma \in \Sigma$ is a (pure) **Nash equilibrium** if it holds

$$C_i(\sigma) \leq C_i(\sigma_{-i}, \sigma'_i) \quad \text{for all } i \in [n], \sigma'_i \in \Sigma_i.$$

Central questions in game theory are the existence of Nash equilibria, their efficiency, and their computational complexity. The efficiency of an equilibrium is measured by the comparison to a social optimum in terms of the social cost. Here, the **social cost** of a strategy profile is defined as $\sum_{i \in [n]} C_i(\sigma)$. A **social optimum** is a minimizer of the social cost. A standard approach to compute Nash equilibria is the **improving dynamics**. Starting with an arbitrary strategy profile, a sequence of improving unilateral deviations is performed. If this process stops, a Nash equilibrium is obtained. In general, however, it is not clear that it terminates.

An example. A famous example for a two player game is the **prisoner's dilemma**, which is attributed to Flood, Dresher, and Tucker (1950), e.g., by Kollock (1998). Two criminals are arrested and interrogated separately. If they keep silence, they can be convicted only on a lesser charge and go to prison for one year each. If both testify, they serve two years in prison each. If one criminal testifies and the other does not, they are set free and the other serves three years. The game is depicted in Table 2.1. The social optimum is for both to keep silence. Then they serve a total of two years. This strategy profile, however, is not a Nash equilibrium. Both individually can reduce their sentence by defecting and testifying. In fact, confessing is a dominant strategy. Thus, the only (pure) Nash equilibrium is the strategy profile in which both confess. Its social cost is a total of four years prison. Therefore, the total time in prison doubles if both criminals testify due to the lack of coordination.

	B keeps silence	B confesses
A keeps silence	1 1	3 0
A confesses	0 3	2 2

Table 2.1: The prisoner’s dilemma. The rows and columns correspond to the strategies of A and B, respectively. The player costs of A and B are colored yellow and violet, respectively.

The efficiency of equilibria. The **price of anarchy** measures the worst-case increase in social cost by the lack of coordination between the players. For an instance of a game, it is defined as the ratio between the maximal social cost of a Nash equilibrium and the social cost of a social optimum. The price of anarchy of a game then is the largest such ratio achieved by any of its instances. This concept is put forward by Koutsoupias and C. Papadimitriou (1999, 2009) under the original name coordination ratio. The widely used term price of anarchy is established by C. H. Papadimitriou (2001).

In some cases, the price of anarchy is too pessimistic as it assumes no central authority at all. A more optimistic view is represented by the **price of stability**. It assumes a central authority that can suggest a strategy profile which, however, needs to be accepted by all players. Formally, the price of stability is defined as the ratio between the minimal social cost of a Nash equilibrium and the social cost of a social optimum. This notion is introduced by Schulz and Stier-Moses (2003) and termed by Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008).

Routing games. In **routing games**, a number of players control the routing of flow through a network. There are two fundamentally different settings. In **atomic** routing games, the flow controlled by the players cannot be split arbitrarily, but only down to some integral unit. In **nonatomic** routing games on the other hand, every infinitesimally small flow particle is considered a player. Hence, in the latter there is a continuum of players. A widely studied solution concept in this setting are **user equilibria** based on Wardrop’s first principle (1952). These are network flows in which flow only travels along shortest paths from its source to its sink. We defer a more formal definition to Chapter 4, where we deal with an extended notion.

Other notions of equilibrium. There are further variations of the concept of an equilibrium. We give three versions of the Nash equilibrium.

A relaxation of (pure) Nash equilibria are **mixed Nash equilibria**. Each player is allowed to randomize over her strategies. A strategy profile is then given by the probability distributions over each player’s strategies. The player costs are determined in expectation with respect to all distributions. In a mixed Nash equilibrium no player can improve their expected cost by randomizing differently. Most remarkably, Nash

(1950) shows that there always exists a mixed Nash equilibrium in every normal form game.

Another relaxation of the Nash equilibrium is an α -**approximate Nash equilibrium** for some $\alpha \geq 1$. It captures the players' inertia. A player deviates only to a different strategy if her cost decreases by at least a factor of α .

A popular refinement of the Nash equilibrium on the other hand is the **strong Nash equilibrium**. In such an equilibrium no coalition of players can decrease the cost of all involved players by cooperatively deviating.

More details on the topic of game theory with a focus on computational aspects can be found in the textbook by Nisan, Roughgarden, Tardos, and Vazirani (2007).

2.5 Computational Problems and Complexity Theory

Complexity theory aims to quantify the amount of resources that is needed to solve various computational problems. Such a **problem** is described by a specific task for a set of possible inputs, called **instances**. The examined resources are typically the number of elementary steps or time, and memory or space. The **Turing machine** serves as the most common abstract computational model and is named after its inventor Alan Turing. In this model, a head manipulates the symbols on an infinitely long tape. The sequence of manipulations is governed by a set of rules, which make up an algorithm. Each step merely depends on the symbol at the head's current position and an internal state. It results in writing a symbol at the current position and moving the head by at most one position. The outcome of a step must be unique for **deterministic Turing machines**. In contrast to that, a **nondeterministic Turing machine** allows multiple outcomes. While a deterministic machine explores a single path of execution, a nondeterministic machine simultaneously explores all possible paths resulting from multiple outcomes. Time is measured as the number of steps a Turing machine takes and space is measured as the number of slots of its infinitely long tape that are read or written. If not stated explicitly, we refer to the computational model of the deterministic Turing machine. Many relevant computational models can be shown to be equivalent to the deterministic Turing machine. This means that the same tasks take essentially the same amount of time and space up to constant factors. Therefore, the complexity of problems is not measured by the exact number of steps and slots needed for solving them. The focus lies more on their asymptotics with respect to the instance sizes. This is where the **Landau symbols** (also called **big O notation**) come in handy. For a function $f: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, the set of functions $\mathcal{O}(f)$ is defined by

$$\mathcal{O}(f) := \{g: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0} \mid \exists N, C > 0 \forall n \geq N: g(n) \leq Cf(n)\}.$$

In other words, $\mathcal{O}(f)$ is the set of functions that do not grow faster than f asymptotically. Similarly, the set of functions growing asymptotically at least as fast as f is defined as

$$\Omega(f) := \{g: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0} \mid \exists N, c > 0 \forall n \geq N: g(n) \geq cf(n)\}.$$

Decision problems. A **decision problem** is stated by a single question on its input that can be answered with yes or no. Formally, a decision problem Π is defined by the set of instances and the subset of instances the answer to which is yes. The inputs are encoded as strings over some fixed alphabet. Typically, binary encoding is chosen for numerical values in the input. The **size** of the input is then the length of its encoding string.

The class \mathbf{P} is the set of all decision problems that can be answered in a number of steps that is polynomial in the size of the input by a deterministic Turing machine. Similarly, the class \mathbf{NP} is the set of decision problems which can be answered in polynomial time by a nondeterministic Turing machine. $\mathbf{P} \subseteq \mathbf{NP}$ follows. The famous question of whether this inclusion holds with equality, that is, whether all problems in \mathbf{NP} can be solved in (deterministic) polynomial time, is one of the seven Millennium Prize Problems stated by the Clay Mathematics Institute (see Jaffe 2006).

Problems can be related to each other by reductions. A **reduction** from a decision problem Π to another decision problem Π' is a mapping from instances of Π to instance of Π' that preserves the answer to the decision problems. If a reduction from Π to Π' can be computed in polynomial time, this shows that Π can be solved in polynomial time whenever Π' can.

This concept allows to identify the hardest problems in \mathbf{NP} . A problem (not necessarily in \mathbf{NP}) is called **NP-hard** if every problem in \mathbf{NP} can be reduced to it in polynomial time. If additionally the problem itself lies in \mathbf{NP} , it is called **NP-complete**. When S. A. Cook (1971) establishes this concept, he shows that the satisfiability problem is NP-complete. Subsequently, Karp (1972) proves the NP-completeness of 21 further problems.

One of these 21 NP-complete problems is **exact 3-set cover**. An instance is given by a ground set U and a family $\mathcal{S} \subseteq \binom{U}{3}$. A **covering** is a subfamily of \mathcal{S} the union of which is U . The goal is to decide whether there exists a covering of size $|U|/3$, that is, a covering which partitions U . A somewhat complementary concept to a covering is the following. A **packing** is a pairwise disjoint subfamily of \mathcal{S} .

Optimization problems. An **optimization problem** asks for the smallest (or largest) objective function value that is attained by a solution in a set of **feasible** solutions. While these are formally no decision problems, they can easily be transformed. The canonical question associated with a minimization problem is whether there is a feasible solution that undercuts a given threshold on the objective function value. If we say an optimization problem is NP-complete, we mean that its canonical decision problem is NP-complete. If a polynomial-time algorithm for this decision variant is given, the optimal value of combinatorial optimization problems can be recovered in polynomial time by binary search (under some mild assumptions). Further, the optimization of combinatorial problems is equivalent to a whole list of related problems as found by Schulz (2009). Note that the definition of optimization problems asks for the optimal value as opposed to a feasible solution that attains this value. It turns out, that an

optimal solution can often be constructed with polynomial overhead if an algorithm for the optimal value is available. This is shown, e.g., for the class of integer programs by Orlin, Punnen, and Schulz (2009).

As many relevant optimization problems are NP-complete, the requirement of optimality is often relaxed. For $\alpha \geq 1$, a feasible solution is an α -**approximation** to a minimization problem if its objective function value is at most α times the objective of an optimal solution. The complexity of computing an α -approximation for fixed α is a central research question for NP-complete optimization problems.

More details on the complexity of decision and optimization problems can be found in the textbook by Garey and Johnson (1990).

Total search problems. Decision problems often ask whether or not a solution exists. Optimization problems on the other hand deal with the existence of solutions with a certain quality measured by a central objective function. Both do not fit problems which are guaranteed to have a solution and the goal is to find an arbitrary one. Those are called **total search problems** or **total function problems**. The class TFNP contains all total search problems a solution to which can be computed by a nondeterministic Turing machine in polynomial time. In contrast to NP, the class TFNP does not seem to have complete problems; see the work by Goldberg and C. H. Papadimitriou (2018a,b). Therefore, several subclasses of TFNP have been established. These classes of problems are typically defined based on a specific method that proves the existence of their solutions.

Johnson, Papadimitriou, and Yannakakis (1985, 1988) define the class **PLS** which contains all problems a solution to which can be found by local search. A problem is defined by an (implicitly defined) set of feasible solutions which are of polynomial size and an objective function that can be evaluated in polynomial time. Further, a set of neighboring feasible solutions is specified for each feasible solution, that can be computed in polynomial time. The goal is to find a **local optimum** which is a feasible solution that has a best objective function value within its **neighborhood**. This class allows the following notion of reduction. A **PLS-reduction** of a problem $\Pi \in \text{PLS}$ to another problem $\Pi' \in \text{PLS}$ is a mapping from instances of Π to instances of Π' with the following properties. Every solution to Π' corresponds to a solution to Π which can be computed in polynomial time. Moreover, every solution to Π that corresponds to a local optimum of Π' must be a local optimum of Π . Under this notion, PLS has **complete** problems that every other problem in PLS reduces to in polynomial time. One such PLS-complete problem, that is relevant to us, is the **maximum cut problem** as shown by Schäffer and Yannakakis (1991). Given an undirected graph $G = (V, E)$ with edge weights $w \in \mathbb{R}_{\geq 0}^E$ the feasible solutions are all cuts $\emptyset \neq U \subsetneq V$. Two cuts U, U' are neighbors if $|U \Delta U'| = 1$. The goal is to find a cut U that locally maximizes the value $w(\delta_E(U))$ with respect to this neighborhood.

Another subclass of TFNP is **PPAD** (short for polynomial parity arguments on directed graphs). It was introduced by C. H. Papadimitriou (1994) and relates to the

proof of existence via fixed-point theorems like the ones by Brouwer (1911) and Kakutani (1941). Formally, it is based on (implicitly defined) directed graphs that are the disjoint union of paths. Given the beginning of a path in such a graph, the end of any path is a solution to the problem. A suitable notion of a reduction can also be defined here. Famously, the work of Daskalakis, Goldberg, and Papadimitriou (2006, 2009) and X. Chen, Deng, and Teng (2009) shows that computing a (mixed) Nash equilibrium in normal form games is PPAD-complete even for two players.

Conditional complexity. While it is conceptually easy to prove that a problem is in \mathbf{P} by providing a polynomial-time algorithm, showing the nonexistence of such an algorithm is challenging in most cases. This is why impossibility statements are often not absolute but relate to a condition which is likely to be true. One such condition, which is widely accepted, is $\mathbf{P} \neq \mathbf{NP}$. It is equivalent to the nonexistence of polynomial-time algorithms for NP-complete problems. Another impossibility statement is the **exponential-time hypothesis** as introduced by Impagliazzo and Paturi (1999, 2001). It states that a variant of the satisfiability problem cannot be solved in subexponential time (in the worst-case).

2.6 Linear Programming

Usually, an optimization problem is formally posed as the task of finding the values for **decision variables** which optimize some **objective function** while satisfying additional **constraints**. An important class of such optimization problems are those with an objective function and (in)equality constraints which are linear in the decision variables. These so-called **linear programs** are capable of modeling many real world applications and are at the same time well-understood theoretically. They are introduced by Kantorovich (1939). Let $n, m \in \mathbb{N}$. The coefficient matrix $A \in \mathbb{R}^{m \times n}$, the right-hand side vector $b \in \mathbb{R}^m$, and the cost vector $c \in \mathbb{R}^n$ define the linear program

$$\max c^\top x \text{ subject to } Ax \leq b, x \geq 0.$$

Here, x is a vector of n decision variables. Note that minimization is equivalent to maximization when replacing the cost vector c with $-c$. A **feasible** solution to the linear program is a vector $x \in \mathbb{R}^n$ which fulfills the constraints $Ax \leq b$ and $x \geq 0$. For index sets $I \subseteq [m]$ and $J \subseteq [n]$ such that $|I| + |J| = n$, the vector

$$\begin{pmatrix} A_{I,\bullet} \\ \text{Id}_{J,\bullet} \end{pmatrix}^{-1} \begin{pmatrix} b_I \\ 0 \end{pmatrix}$$

is called a **basic** solution, if the involved matrix is invertible. The maximal value $c^\top x$ that is achieved by any feasible solution x is called the **value** val of the linear program. The solution x achieving it is called **optimum** or **optimal** solution. A linear program is **bounded** if its value is finite.

Duality. There is an intriguing relation between pairs of linear programs, called **duality**. The **dual linear program** of the preceding (primal) linear program is given by

$$\min b^\top y \text{ subject to } A^\top y \geq c, y \geq 0.$$

Note that this relation is symmetric. The primal linear program is the dual of its dual. Every variable of the primal is associated with a constraint in the dual, every constraint of the primal corresponds to a variable of the dual. The feasible solutions $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ to the primal and dual linear program, respectively, are optimal if and only if they fulfill **complementary slackness**

$$y^\top (Ax - b) = 0 \quad \text{and} \quad x^\top (A^\top y - c) = 0.$$

This immediately implies the **strong duality theorem**. If $x^* \in \mathbb{R}^n$ and $y^* \in \mathbb{R}^m$ are optimal solutions to the primal and dual linear program, respectively, it holds

$$c^\top x^* = (y^*)^\top Ax^* = b^\top y^*.$$

In particular, the objective value of any feasible primal solution is larger than the objective value of any feasible dual solution. This is known as the **weak duality theorem**.

Polyhedra and polytopes. A subset $P \subseteq \mathbb{R}^n$ is called **polyhedron** if there is a matrix $A \in \mathbb{R}^{m \times n}$ and a vector $b \in \mathbb{R}^m$ such that $P = \{x \in \mathbb{R}^n \mid Ax \leq b\}$. A bounded polyhedron is called **polytope**. For $I \subseteq [m]$, the subset $\{x \in P \mid A_{I,\bullet} x = b_I\}$ is a **face** of P . The **dimension** of a face $F \neq \emptyset$ is the dimension of the linear subspace $F - x$ with $x \in F$. A face of dimension zero or one is called **vertex** and **edge**, respectively.

Clearly, the set of feasible solutions to a linear program is a polyhedron. Its set of vertices corresponds to the set of basic feasible solutions. If a linear program is bounded and its feasible set has at least one vertex, there exists a vertex which is an optimal (basic) solution.

Computational complexity. There are several methods to solve linear programs algorithmically. Khachiyan (1979) applies the **ellipsoid method** to linear programming and thereby finds the first method to solve linear programs in polynomial time. Its implementation, however, poses substantial numerical issues, which is why it is rarely used in practice. The approach which enjoys the greatest popularity is the **simplex algorithm**. It was developed by Dantzig in 1947 (see Dantzig 1990) and is basically a local search on the vertices of the feasible polyhedron. Two vertices are considered neighbors if they are contained in a common edge of the polyhedron. Switching between two such neighboring vertices is also called **pivoting** based on the underlying algebra. After finding an arbitrary vertex, the algorithm follows an improving path of neighboring vertices until no further improvement can be achieved in this way. The initial feasible vertex is found by a similar process, called **phase I**. It applies a local search to a modified version of the linear program which has a trivial feasible solution. Despite the fact that Klee and Minty (1972) show an exponential worst-case running time, the algorithm achieves high

efficiency in practice. A theoretical explanation thereof is provided by Spielman and Teng (2001, 2004). They essentially show that the simplex algorithm runs in polynomial time for most inputs.

Integer linear programs. In many applications, fractional values for decision variables do not have a meaningful interpretation. Therefore, the constraint for variables to be integral arises naturally. A linear program with the additional constraint that all its variables are integral is called **integer linear program**. If both continuous and integral variables are present, it is called **mixed-integer linear program**. Both these problems are NP-complete. Some problems allow structural insights that lead to efficient combinatorial algorithms. General algorithms are typically based on linear programming with additional techniques to establish integrality, e.g., **branch and bound** or **cutting plane** methods.

An example. We present a famous duality result from network flows. A formulation for the **maximum s - t flow problem** on a directed graph $G = (V, A)$ with capacities $\nu \in \mathbb{R}_{\geq 0}^A$, source $s \in V$, and sink $t \in V$ is given by

$$\begin{aligned} \max \quad & \sum_{a \in \delta_A^-(t)} x_a - \sum_{a \in \delta_A^+(t)} x_a && \text{(MAXFLOW)} \\ \text{s. t.} \quad & \sum_{a \in \delta_A^-(v)} x_a - \sum_{a \in \delta_A^+(v)} x_a \geq 0 && \text{for all } v \in V \setminus \{s, t\} \\ & 0 \leq x_a \leq \nu_a && \text{for all } a \in A. \end{aligned}$$

On every arc $a \in A$ the amount of flow is described by the nonnegative decision variable x_a . The capacity constraints bound the flow value x_a on an arc $a \in A$ by its capacity ν_a . Further, weak flow conservation is imposed on all nonterminal vertices. The objective is to maximize the value of the flow. Note that by the weak flow conservation excess at the sink t can only emerge from the source s . Requiring strong instead of weak flow conservation for the maximum flow problem yields an equivalent linear program. This is because flow that leaves the source but never reaches the sink can be discarded without making the flow infeasible or changing the objective function value.

The dual linear program of (MAXFLOW) is

$$\begin{aligned} \min \quad & \sum_{a \in A} \nu_a \mu_a && \text{(MINCUT)} \\ \text{s. t.} \quad & \mu_a \geq \lambda_w - \lambda_v && \text{for all } a = (v, w) \in A \\ & \mu_a \geq 0 && \text{for all } a \in A \\ & \lambda_v \geq 0 && \text{for all } v \in V \setminus \{s, t\}, \quad \lambda_s = 0, \quad \lambda_t = 1 \end{aligned}$$

We can assume that an optimal solution fulfills $\mu_a = [\lambda_w - \lambda_v]_+$ for all $a = (v, w) \in A$. Further, it can be shown that there is always an optimal integral solution $\lambda \in \{0, 1\}^A$ to (MINCUT). Its objective function value is exactly the capacity of the (directed) cut $\delta_A^+(\{v \in V \mid \lambda_v = 0\})$. Strong duality now yields that the maximum value of an s - t flow is equal to the minimum value of an s - t cut which is a cut $U \subseteq V$ with $s \in U$ and $t \notin U$. This is the famous **max-flow min-cut theorem** as found by Elias, Feinstein, and Shannon (1956) and Ford and Fulkerson (1956). Under the assumption that the capacities ν are integral, it can be shown that the feasible set of (MAXFLOW) has integral vertices only. Together with the observation, that the program is always bounded if $s \neq t$ it follows that there always exists an optimal integral flow under this assumption.

For more on the topic of linear and integer programming, we recommend the textbook by Schrijver (1998).

2.7 Linear Complementarity Problems

Complementarity conditions appear naturally in optimality criteria for linear programming and other optimization problems, like quadratic programs. Further, they arise from the conditions on a Nash equilibrium. This motivates the following class of feasibility problems. A **linear complementarity problem** is defined by a square matrix $M \in \mathbb{R}^{n \times n}$ and a vector $q \in \mathbb{R}^n$. The goal is to find a vector $z \in \mathbb{R}^n$ that satisfies

$$z^\top (Mz + q) = 0, \quad Mz + q \geq 0, \quad z \geq 0.$$

By introducing additional auxiliary variables $w \in \mathbb{R}^n$, the problem can be restated in a more symmetric fashion. It is equivalent to finding w and z such that

$$w = Mz + q, \quad z^\top w = 0, \quad w \geq 0, \quad z \geq 0.$$

The condition $z^\top w = 0$ is called **complementarity condition**. Under the assumption that w and z are nonnegative, it is equivalent to requiring $w_i = 0$ or $z_i = 0$ for all $i \in [n]$. In terms of the preceding formulation, $M_{i,\bullet} z + q_i \geq 0$ or $z_i \geq 0$ has to be fulfilled with equality for every $i \in [n]$.

Computational complexity. The class of linear complementarity problems is quite heterogeneous in terms of computational complexity. Not surprisingly, the decision problem on the existence of a solution is NP-complete. Consequently, subsets of problems are treated which are guaranteed to have a solution. These problem classes are typically defined by properties of the matrix M . An example of linear complementarity problems that always have a unique solution are those with symmetric positive definite matrix M . Restricting to such a class yields a total search problem.

The known algorithms for solving linear complementarity problems can be grouped in pivoting and iterative methods. Their applicability and efficiency generally depends on the matrix M . A popular pivoting algorithm is **Lemke's algorithm** (1965). It starts

by relaxing the given linear complementary problem by the introduction of an auxiliary variable. Pivoting steps then follow a path of solutions that fulfill all but one of the original complementarity conditions. If the auxiliary variable can be finally eliminated, a solution is found. We defer more specific results on Lemke's scheme to [Chapter 4](#) where they are used.

An example. An important class of linear complementarity problems which contributed to their introduction is given by two-player normal form games, also termed **bimatrix games**. Let $\Sigma = \Sigma_1 \times \Sigma_2$ be the strategy profiles of such a game, and let $A, B \in \mathbb{R}^\Sigma$ be two matrices that describe the player costs. More specifically, A_{σ_1, σ_2} and B_{σ_1, σ_2} are the player costs of player A and B under the profile $(\sigma_1, \sigma_2) \in \Sigma$, respectively. We may assume that the entries of A and B are positive, because adding the same scalar to all entries results in an equivalent game. Then the solutions x, y to the linear complementarity problem

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} \mathbb{1} \\ \mathbb{1} \end{pmatrix}, \quad \begin{pmatrix} u \\ v \end{pmatrix}^\top \begin{pmatrix} x \\ y \end{pmatrix} = 0, \quad u, v \geq 0, \quad x, y \geq 0 \quad (\text{BIMATRIX})$$

correspond one-to-one to the mixed Nash equilibria via $(x, y) \mapsto (x/\mathbb{1}^\top x, y/\mathbb{1}^\top y)$. More precisely, $x_{\sigma_1}/\mathbb{1}^\top x$ is the probability that player A plays the strategy $\sigma_1 \in \Sigma_1$, and similarly for player B. The mapping normalizes x and y such that we indeed obtain probability distributions. On the one hand, $Ay \geq \mathbb{1}$ holds. On the other hand, the complementarity constraints imply $A_{\sigma_1, \bullet} y = 1$ or $x_{\sigma_1} = 0$ for every $\sigma_1 \in \Sigma_1$. In other words, player A can only play a strategy σ_1 with positive probability if it provides minimal expected cost $A_{\sigma_1, \bullet} y/\mathbb{1}^\top y$ under player B's mixed strategy $y/\mathbb{1}^\top y$. The analogous property is obtained for player B. Hence, a solution to (BIMATRIX) indeed corresponds to a Nash equilibrium. The inverse can be seen when scaling the probability distributions of a mixed Nash equilibrium such that the expected costs become one.

A pivoting algorithm, now known as Lemke-Howson algorithm, tailored specifically to bimatrix games is found by Lemke and Howson ([1964](#)).

For more details on linear complementarity problems, see the textbook of Cottle, Pang, and Stone ([2009](#)).

Chapter 3

Two Disjoint Shortest Paths

In this chapter, we consider a specific multicommodity integer flow problem. The edges of the underlying undirected networks can carry at most one unit of flow, which excludes multiple commodities using the same edge. Each commodity needs to ship exactly one unit along a path. Additionally, this path is required to be a shortest path between the given source and sink. For k commodities this problem is known as the **(undirected) k disjoint shortest paths problem**. In other words, it asks for the existence of k pairwise edge- or vertex-disjoint shortest s_i - t_i paths in an undirected graph with k source-sink pairs $(s_i, t_i), i \in [k]$. If k is not fixed but part of the input, the problem is NP-complete. The **two disjoint shortest paths problem** ($k = 2$) with strictly positive edge lengths on the other hand can be solved in polynomial time. Our main contribution is the generalization of this tractability to **nonnegative edge lengths**. Structural insights allow us to deal with connected components of zero-length edges separately from edges with strictly positive length. The resulting algorithm is based on dynamic programming and runs in polynomial time.

Authorship. The findings of this chapter are a result of a collaboration with Marinus Gottschau and Clara Waldmann. The presentation has similarities with the article of Gottschau, Kaiser, and Waldmann (2019) where the results have been published previously. Our main insight has been obtained independently by Kobayashi and Sako (2019).

Outline. Section 3.1 starts out by introducing the problem. It then relates our work to existing literature and states our main results on a high level. The body of this chapter is organized into three major parts. Section 3.2 builds the foundation for the following two sections. A decomposition of graphs with an acyclic structure is developed, which allows to split the task of finding disjoint paths into smaller parts. In Section 3.3, we present an intermediate result. More specifically, we treat k disjoint paths in weakly acyclic mixed graphs, which generalize directed acyclic graphs. We obtain a polynomial-time dynamic program based on the previously established decomposition. In Section 3.4, this intermediate result is combined with another application of the decomposition. The outcome is a first polynomial-time algorithm for the two disjoint shortest paths problem with nonnegative edge lengths. Section 3.5 closes with some remarks on possible research directions.

3.1 Introduction

We start by introducing the problem at hand formally. Before doing so, we define a more basic variant, the disjoint paths problem.

Disjoint paths. The k **disjoint paths problem** represents a very fundamental and well-studied problem in graph theory. For a given graph G and k source-sink pairs $(s_i, t_i), i \in [k]$, it asks for the existence of k pairwise disjoint paths P_1, \dots, P_k such that P_i is an s_i - t_i path for all $i \in [k]$. Versions of the problem result from considering either disjointness of the paths with respect to vertices or edges. Further, the underlying graph can be directed or undirected. Unless stated otherwise, we refer to an undirected graph $G = (V, E)$.

Disjoint shortest paths. The problem we treat in this chapter can be viewed as generalization of the disjoint paths problem. In the k **disjoint shortest paths problem**, additional requirements on the paths $P_i, i \in [k]$ are introduced by edge lengths $\ell \in \mathbb{R}_{\geq 0}^E$. For every $i \in [k]$, the path P_i needs to be a shortest s_i - t_i path with respect to ℓ . Note that the k disjoint paths problem is recovered when setting $\ell \equiv 0$. Hence, we indeed deal with a generalization. The discussed versions of the problem generalize accordingly. Furthermore, the variants with unit edge lengths ($\ell \equiv 1$) take a special role.

For many disjoint paths problems, there is a generic reduction from finding arc/edge-disjoint paths to finding vertex-disjoint paths. This reduction is based on the **line graph**. Given a graph G , its line graph has a vertex for every arc and edge in G . There is an arc or edge between two vertices in the line graph if the corresponding arcs/edges are incident to a common vertex in the original graph G (and their directions are compatible); see [Figure 3.1](#). Then arc/edge-disjoint paths in G correspond to vertex-disjoint paths in its line graph. For directed graphs, there is also an opposite reduction from vertex-disjoint to arc-disjoint variants based on splitting vertices; see [Figure 3.2](#). This generic reduction, however, cannot be applied to undirected or mixed graphs.

Consequently, the vertex-disjoint and arc-disjoint variants are equivalent when considering directed disjoint paths. For undirected and mixed graphs, the vertex-disjoint variant is generally at least as hard as the corresponding arc/edge-disjoint problem. We focus on the arc/edge-disjoint variants and discuss the vertex-disjoint equivalents subsequently.

3.1.1 Related Literature

We give a brief overview of the most important results related to the work presented in this chapter. We start with the more basic disjoint paths problem and continue with the disjoint shortest paths problem and a generalization, the shortest disjoint paths problem. Finally, we highlight some results in the allied field of integer flow. An overview of the complexity results on disjoint paths problems is given in [Table 3.1](#).



Figure 3.1: The construction of a line graph for mixed graphs. Every arc or edge in the original graph (left) is associated with a vertex in the line graph (right). Vertices representing two incident edges are connected by an edge. Vertices representing two incident arcs/edges which form a path are connected by an arc.



Figure 3.2: A reduction from directed vertex-disjoint paths to directed arc-disjoint paths. Every vertex of the original graph (left) is split into two vertices connected by a new arc (right). The incoming and outgoing arcs are separated by this splitting.

Disjoint paths. The general case of the k disjoint paths problem is NP-hard when k is considered part of the input. This is obtained by Karp (1975) for vertex-disjoint paths and by Even, Itai, and Shamir (1975, 1976) for edge-disjoint paths. It remains hard when restricting to planar graphs as found by Lynch (1975). Therefore, a lot of research focuses on the setting where k is considered to be fixed.

For the undirected two disjoint paths problem, polynomial-time algorithms were obtained by Seymour (1980), Shiloach (1980), and Thomassen (1980). In their seminal article, Robertson and Seymour (1995) provide an algorithm for the undirected k disjoint paths problem with arbitrary fixed k that runs in $\mathcal{O}(|V|^3)$. Kawarabayashi, Kobayashi, and Reed (2012) reduce this running time to $\mathcal{O}(|V|^2)$.

In contrast to that, the directed version of the problem turns out to be much harder. Fortune, Hopcroft, and Wyllie (1980) prove that even the directed two disjoint paths problem is NP-hard. For directed acyclic graphs and arbitrary fixed k the authors find a polynomial-time algorithm. Zhang and Nagamochi (2012) extend this work to solve the problem on acyclic mixed graphs. Here, a mixed graphs is called acyclic if orienting any set of edges does not produce a directed cycle.

Disjoint shortest paths. Eilam-Tzoref (1998) introduces the disjoint shortest paths problem as a generalization of the disjoint paths problem. She shows that the problem is NP-hard if k is part of the input even in planar graphs with unit-length edges ($\ell \equiv 1$). If zero-length edges are allowed, the hardness of the k disjoint shortest paths problem is already implied by the hardness of the k disjoint paths problem. In particular, the directed two disjoint shortest paths problem is NP-hard. For both the undirected and

	undirected			directed		
	$\ell \equiv 0$	$\ell \geq 0$	$\ell > 0$	$\ell \equiv 1$	$\ell \equiv 0$	$\ell > 0$
k arbitrary	NP-hard ¹	NP-hard ¹	NP-hard ²	NP-hard ²	NP-hard ¹	NP-hard ²
k fixed	P ³			P ⁴	NP-hard ⁵	
$k = 2$	P ⁶	P ⁷	P ²	P ²	NP-hard ⁵	P ⁸

Table 3.1: The complexity of the disjoint paths problem ($\ell \equiv 0$) and its variants.

directed setting, the problem with $k = 2$ becomes solvable in polynomial time when restricting to strictly positive edge lengths. For undirected graphs, Eilam-Tzoreff (1998) provides a dynamic program. Directed graphs are treated by Bérczi and Kobayashi (2017). In contrast to the directed variant, the undirected variant for $k = 2$ can still be solved efficiently for nonnegative edge lengths, as we see in this chapter. This result is obtained independently by Kobayashi and Sako (2019).

Lochet (2021) shows that the undirected k disjoint shortest paths problem can be solved in polynomial time for arbitrary fixed k if all edges have unit length ($\ell \equiv 1$). His algorithm runs in $\mathcal{O}(|V|^{k^{5^k}})$. Further, the author shows that the problem is $W[1]$ -hard with respect to the parameter k . Under the exponential time hypothesis, this implies that there is no algorithm that runs in time $\mathcal{O}(f(k) \cdot |V|^n)$ for any fixed $n \in \mathbb{N}$ and function $f: \mathbb{N} \rightarrow \mathbb{N}$. In other words, the dependency of the exponent on k is most likely inevitable. Bentert, Nichterlein, Renken, and Zschoche (2021) reduce this dependency to $\mathcal{O}(k \cdot |V|^{12kk!+k+1})$.

For arbitrary fixed k but under the assumption that the graph is planar, Bérczi and Kobayashi (2017) achieve polynomial-time solvability for the undirected cases and the directed k vertex-disjoint shortest paths problem.

Shortest disjoint paths. A further generalization of the disjoint shortest paths problem is the **shortest disjoint paths problem**. Instead of requiring that every single path is shortest, a family of disjoint paths is searched which minimizes the total length. In an early work, Suurballe (1974) uses augmentation techniques to devise a polynomial-time algorithm minimizing the total length, if all paths share a common source and sink. Björklund and Husfeldt (2014, 2019) find a polynomial-time algebraic Monte Carlo algorithm for solving the shortest disjoint paths problem with unit lengths ($\ell \equiv 1$).

¹Even, Itai, and Shamir (1975, 1976) and Karp (1975)

²Eilam-Tzoreff (1998)

³Robertson and Seymour (1995), improved by Kawarabayashi, Kobayashi, and Reed (2012)

⁴Lochet (2021), improved by Bentert, Nichterlein, Renken, and Zschoche (2021)

⁵Fortune, Hopcroft, and Wyllie (1980)

⁶independently by Seymour (1980), Shiloach (1980), and Thomassen (1980)

⁷our main result, independently found by Kobayashi and Sako (2019)

⁸Bérczi and Kobayashi (2017)

Multicommodity integer flow. The multicommodity integer flow problem can be viewed as a generalization of the disjoint paths problem. Indeed, the first proofs of NP-hardness of the disjoint paths problem by Karp (1975) and Even, Itai, and Shamir (1975, 1976) are stated in terms of flows. The latter authors show that deciding on the existence of multicommodity integer flows is already NP-hard for two commodities.

On the positive side, there are some results of (singlecommodity) integer flows which transfer to special settings of disjoint paths problems. Menger’s theorem (1927) represents an early result on a special case of strong duality. For vertices $s, t \in V$, it states that the size of a minimum s - t cut agrees with the maximum number of pairwise edge-disjoint s - t paths. This is generalized to the famous max-flow min-cut theorem found by Elias, Feinstein, and Shannon (1956) and Ford and Fulkerson (1956). Introducing capacities, it shows that the minimum capacity of an s - t cut is the same as the maximum value of an s - t flow. Consequently, known algorithms for computing a maximum s - t flow yield that the k disjoint paths problem can be decided in polynomial time if all sources and sinks agree, even if k is part of the input (e.g., the improved algorithms of Ford and Fulkerson (1956) by Dinic (1970) and Edmonds and Karp (1972)). By standard techniques, this extends to the cases where only the sources or the sinks are the same. Similarly, the minimum-cost integer flow problem is a generalization of the shortest disjoint paths problem if all sources or sinks are the same.

3.1.2 Our Contribution

We give a first polynomial-time algorithm for the undirected two disjoint shortest paths problem with nonnegative edge lengths. In order to deal with zero-length edges, we combine techniques by Fortune, Hopcroft, and Wyllie (1980) and Bérczi and Kobayashi (2017). The two shortest paths problem is transformed into a disjoint paths problem without lengths. The requirement of the paths to be shortest paths is implemented by restricting them to their respective shortest path network. We then decompose the transformed problem into the parts where the two shortest path networks overlap and the remainder. The overlapping parts exhibit a specific structure which we call **weakly acyclic**. It allows a further decomposition separating edges with zero and positive lengths.

3.2 A Decomposition of Disjoint Paths

To express the existence of disjoint paths concisely, we introduce the following notation.

Definition 3.1 (Disjoint paths relations). Let $G = (V, \mathcal{E})$ be a (mixed) graph with nonnegative edge lengths $\ell \in \mathbb{R}_{\geq 0}^{\mathcal{E}}$. For $k \in \mathbb{N}$, we define the two binary relations $\rightrightarrows_{\mathcal{E}}$ and $\xrightarrow{\ell}_{\mathcal{E}}$ on the set V^k .

- (i) For $v, w \in V^k$, we set $v \rightrightarrows_{\mathcal{E}} w$ if there exist pairwise arc/edge-disjoint paths P_1, \dots, P_k such that P_i is a v_i - w_i path in \mathcal{E} for all $i \in [k]$. We also use \rightrightarrows_G for $\rightrightarrows_{\mathcal{E}}$.

- (ii) For $v, w \in V^k$, we set $v \xrightarrow{\ell} w$ if there exist pairwise arc/edge-disjoint paths P_1, \dots, P_k such that P_i is a shortest v_i - w_i path with respect to ℓ for $i \in [k]$. We also write $\xrightarrow{\ell}_G$ instead of $\xrightarrow{\ell}_E$.

Note that for fixed k the relation $\Rightarrow_{\mathcal{A}}$ has polynomial size in $|V|$. Our algorithms compute large parts of the disjoint paths relations. Hence, they decide the existence of disjoint paths for multiple source and sink tuples simultaneously. For a single tuple of sources and sinks, improvements in the running time might be achieved.

We show that the disjoint paths relation can be decomposed into smaller relations if the underlying graph exhibits a certain acyclic structure. We start by observing basic properties of the disjoint paths relations. The relations $\Rightarrow_{\mathcal{A}}$ and $\xrightarrow{\ell}_E$ are reflexive; that is, $v \Rightarrow_{\mathcal{A}} v$ and $v \xrightarrow{\ell}_E v$ hold for all $v \in V^k$. This is due to the fact that a collection of empty paths is trivially pairwise arc/edge-disjoint (and shortest). While $\Rightarrow_{\mathcal{A}}$ is generally not transitive, we still observe a similar property between two disjoint paths relations on two disjoint sets of arcs and edges. Let \mathcal{A}_1 and \mathcal{A}_2 be such two disjoint sets of arcs and edges. Let $u, v, w \in V^k$ be tuples of vertices such that $u \Rightarrow_{\mathcal{A}_1} v$ and also $v \Rightarrow_{\mathcal{A}_2} w$. Then we obtain that $u \Rightarrow_{\mathcal{A}_1 \cup \mathcal{A}_2} w$ holds. This is simply based on the concatenation of the disjoint paths from u to v in \mathcal{A}_1 with the respective disjoint paths from v to w in \mathcal{A}_2 . In other words, we have the inclusion of relations

$$\Rightarrow_{\mathcal{A}_2} \circ \Rightarrow_{\mathcal{A}_1} \subseteq \Rightarrow_{\mathcal{A}_1 \cup \mathcal{A}_2}.$$

The following theorem gives sufficient conditions for this inclusion to hold with equality.

Theorem 3.2 (Decomposition of the disjoint paths relation). *Let $G = (V, A \cup E)$ be a mixed graph. Further, let $V = \bigcup_{j=1}^h V_j$ be a partition such that the contraction $G/\{V_1, \dots, V_h\}$ yields a directed acyclic graph. Assume that V_1, \dots, V_h are indexed in topological order with respect to it. Then the disjoint paths relation \Rightarrow_G decomposes into*

$$\Rightarrow_{G[V_h]} \circ \Rightarrow_{\delta_A^-(V_h)} \circ \Rightarrow_{G[V_{h-1}]} \circ \dots \circ \Rightarrow_{\delta_A^-(V_3)} \circ \Rightarrow_{G[V_2]} \circ \Rightarrow_{\delta_A^-(V_2)} \circ \Rightarrow_{G[V_1]}.$$


Proof. The set $A \cup E$ of arcs and edges is partitioned by

$$\mathcal{A}(G[V_1]) \cup \delta_A^-(V_2) \cup \mathcal{A}(G[V_2]) \cup \delta_A^-(V_3) \cup \dots \cup \mathcal{A}(G[V_{h-1}]) \cup \delta_A^-(V_h) \cup \mathcal{A}(G[V_h]).$$


Based on the discussion preceding the theorem, all tuples related by the composition of the disjoint paths relations on these sets are also related by \Rightarrow_G .

For the opposite inclusion, let $v, w \in V^k$ such that $v \Rightarrow_G w$. Hence, there exist pairwise arc/edge-disjoint paths P_1, \dots, P_k in G such that P_i is a v_i - w_i path for all $i \in [k]$. For every $i \in [k]$, we may assume without loss generality that P_i is simple, that is, it does not visit any vertex twice. (Otherwise, we could shortcut.) We find a suitable partition of the paths by intermediate tuples of vertices to show v and w are also related in the composed relation. For every $j \in [h]$, define the vertex tuples $p^{(j)}, q^{(j)} \in V^k$ as follows. For $i \in [k]$ and $j \in [h]$, let $p_i^{(j)}$ be the first vertex on P_i that is in the set $\bigcup_{l=j}^h V_l$. If such


3.2 A Decomposition of Disjoint Paths



(a) A mixed graph and a partition of its vertices.



(b) The directed acyclic graph resulting from contracting the partition.



(c) The decomposition of two disjoint paths by intermediate vertices based on the graph's partition.

Figure 3.3: A graph and a partition which allows the decomposition in the sense of [Theorem 3.2](#).

a vertex does not exist, define $p_i^{(j)}$ as w_i instead. Further, set $q_i^{(j)}$ to be the last vertex on P_i from V_j . If it does not exist, set $q_i^{(j)}$ equal to $p_i^{(j)}$. See [Figure 3.3](#) for an illustration.

We argue that these intermediate tuples fulfill

$$v = p^{(1)} \rightrightarrows_{G[V_1]} q^{(1)} \rightrightarrows_{\delta_A^-(V_2)} p^{(2)} \rightrightarrows_{G[V_2]} q^{(2)} \rightrightarrows_{\delta_A^-(V_3)} \cdots \rightrightarrows_{\delta_A^-(V_h)} p^{(h)} \rightrightarrows_{G[V_h]} q^{(h)} = w.$$

As $G/\{V_1, \dots, V_h\}$ is acyclic and V_1, \dots, V_h are ordered topologically, we have that

$$v_i = p_i^{(1)}, q_i^{(1)}, p_i^{(2)}, q_i^{(2)}, \dots, p_i^{(h)}, q_i^{(h)} = w_i$$

appear on P_i in that order for all $i \in [k]$. If P_i visits V_j , then the part of P_i from $p_i^{(j)}$ to $q_i^{(j)}$ completely lies within $G[V_j]$. Otherwise, $p_i^{(j)} = q_i^{(j)}$ holds. In any case, we obtain that $p^{(j)} \rightrightarrows_{G[V_j]} q^{(j)}$ because the subpaths $P_i \cap G[V_i]$ are pairwise arc/edge-disjoint $p_i^{(j)}-q_i^{(j)}$ paths for all $i \in [k]$. To see $q^{(j)} \rightrightarrows_{\delta_A^-(V_{j+1})} p^{(j+1)}$ for $j \in [h-1]$, we consider a similar case distinction. If P_i visits V_j and $q_i^{(j)} \neq w_i$, then $q_i^{(j)}$ is the last vertex on P_i within V_j , and $p_i^{(j+1)}$ is the first vertex on P_i that lies in $\bigcup_{l=j+1}^h V_l$. Thus, there must be a single arc in P_i from $q_i^{(j)}$ to $p_i^{(j+1)}$. In the remaining case, we have $q_i^{(j)} = p_i^{(j+1)}$. In both cases, the necessary relation holds. \blacksquare

Algorithm 1: Dynamic program for the k disjoint paths problem in weakly acyclic mixed graphs

Input: weakly acyclic mixed graph $G = (V, A \cup E)$

Output: \Rightarrow_G on V^k

- 1 Find connected components V_1, \dots, V_h of the subgraph (V, E) in topological ordering with respect to $G/\{V_1, \dots, V_h\}$;
 - 2 **for** $j = 1, \dots, h$ **do** compute $\Rightarrow_{G[V_j]}$ and $\Rightarrow_{\delta_A^-(V_j)}$;
 - 3 Initialize \Rightarrow to be the identity relation $\{(v, v) \mid v \in V^k\}$;
 - 4 **for** $j = 1, \dots, h$ **do** update \Rightarrow to $\Rightarrow_{G[V_j]} \circ \Rightarrow_{\delta_A^-(V_j)} \circ \Rightarrow$;
 - 5 **return** \Rightarrow ;
-

3.3 Disjoint Paths in Weakly Acyclic Mixed Graphs

This section treats an intermediate problem, which might be interesting on its own. We examine disjoint paths in what we call **weakly acyclic mixed graphs**. These are mixed graphs that are directed acyclic graphs with respect to their undirected connected components. This structure allows to separate the treatment of the directed and undirected parts. It appears naturally when orienting shortest path networks in the [Section 3.4](#).

Definition 3.3 (Weakly acyclic mixed graphs). *A mixed graph $G = (V, A \cup E)$ is **weakly acyclic** if the contraction of its connected components with respect to E does not contract any arcs in A and yields a directed acyclic graph.*

It can be seen that a mixed graph is weakly acyclic if and only if it does not contain a cycle with at least one arc from A . Note that a weakly acyclic mixed graph can indeed have (undirected) cycles in its edge set E .

We propose [Algorithm 1](#) to compute k disjoint paths in weakly acyclic mixed graphs. The fundamental idea of our approach is to treat edges and arcs separately. The set of edges is further partitioned into its connected components. Let V_1, \dots, V_h be the connected components of (V, E) . By definition, their contraction yields a directed acyclic graph $G/\{V_1, \dots, V_h\}$. We assume that V_1, \dots, V_h are in topological order with respect to it. As there cannot be any directed arcs within a connected component V_i , the set of edges is partitioned by the undirected induced subgraphs $G[V_1], \dots, G[V_h]$. Further, we obtain a partition of the set of arcs into $\delta_A^-(V_1), \dots, \delta_A^-(V_h)$. Due to [Theorem 3.2](#), the disjoint paths relation on G can be found by composing the relations on

$$G[V_1], \delta_A^-(V_2), G[V_2], \delta_A^-(V_3), \dots, G[V_{h-1}], \delta_A^-(V_h), G[V_h].$$




Figure 3.4: In iteration j of [Algorithm 1](#), the relation \Rightarrow is extended by concatenating previously computed paths with pairwise different arcs from $\delta_A^-(V_j)$ and edge-disjoint paths in $G[V_j]$.

For the undirected connected components, a subroutine for computing edge-disjoint paths in undirected graphs is used to obtain the disjoint paths relation, e.g., the algorithm by Kawarabayashi et al. (2012). Note that the directed parts between the undirected components have a very simple structure. Therein, all paths are arc-disjoint. The composition of these relations is obtained via dynamic programming. The disjoint paths relation \Rightarrow is computed on successively larger parts of the mixed graph. Found arc/edge-disjoint paths are extended alternately by arcs between components and edge-disjoint paths within one component; see [Figure 3.4](#). This approach represents an extension of the work by Fortune et al. (1980). In the case $E = \emptyset$, the undirected connected components are all single vertices and our algorithm agrees with theirs for directed acyclic graphs.

Theorem 3.4 (Correctness and complexity). *Let $k \in \mathbb{N}$ be fixed. Given a weakly acyclic mixed graph $G = (V, A \cup E)$, [Algorithm 1](#) computes the disjoint paths relation \Rightarrow_G on V^k in polynomial time.*

Proof. Let $V = \bigcup_{j=1}^h V_j$ be the partition of V into the vertex sets of the h connected components of (V, E) as computed by the algorithm.

Correctness. For all $j \in \{0, \dots, h\}$, let \mathcal{A}_j be the arc and edge set of $G[\bigcup_{i=1}^j V_i]$. In particular, $\mathcal{A}_0 = \emptyset$ holds. It follows by induction from [Theorem 3.2](#) that the computed relation \Rightarrow is equal to $\Rightarrow_{\mathcal{A}_j}$ after the j -th iteration of Line 4. In particular, the returned relation is indeed the disjoint paths relation on G .

Complexity. The connected components of $V_j, j \in [h]$ and their topological ordering with respect to $G/\{V_1, \dots, V_h\}$ can be computed in polynomial time, e.g., using Kahn's algorithm (1962). As already discussed, edge-disjoint paths in the undirected components can also be found efficiently. Hence, the relation $\Rightarrow_{G[V_j]}$ can be computed in polynomial time by full enumeration of all possible pairs of source and sink vectors. The relations $\Rightarrow_{\delta_A^-(V_j)}$ can be obtained by enumerating (ordered) subsets of arcs. A binary relation on V^k contains at most $|V|^{2k}$ elements. Therefore, the composition of polynomially many such relations can be computed in polynomial time. In total, we see that [Algorithm 1](#) runs in time polynomial in the size of the input if k is fixed. ■

As discussed in [Section 3.1](#), the vertex-disjoint and arc/edge-disjoint variants are generally not equivalent in mixed graphs. Still, [Algorithm 1](#) can be adapted to compute the relation for vertex-disjoint paths. First, the disjoint paths relations for the parts of the decomposed graph have to be computed in a vertex-disjoint sense. Then, they are defined on the ground set $\{v \in V^k \mid v_1, \dots, v_k \text{ pairwise distinct}\}$. For the undirected connected components, a suitable algorithm for undirected vertex-disjoint paths is used. For the directed parts in between, the disjoint paths relation is given by all possible matchings of size at most k .

3.4 Undirected Disjoint Shortest Paths

In this section, we present the main result of this chapter. We derive a polynomial-time algorithm for the two disjoint shortest paths in undirected graphs with nonnegative edge lengths. Our approach is an extension of the techniques introduced by Bérczi and Kobayashi ([2017](#)). By restricting to shortest path networks, we transform the undirected graph to a mixed graph on which we solve a disjoint paths problem. The challenge of dealing with edges of length zero is tackled by applying the results of [Section 3.3](#).

3.4.1 Orienting Shortest Paths

Given an instance of the two disjoint shortest paths problem, we construct an equivalent variant of a two disjoint paths problem on a mixed graph. Essentially, the requirement of using shortest paths is implemented by restricting to shortest path networks and subsequently dropping the edge lengths. As the two source-sink pairs adhere to different shortest path networks, we need to make this distinction also in the transformed graph.

Let a two disjoint shortest paths problem be given by an undirected graph $G = (V, E)$, source and sink vectors $s, t \in V^2$, and nonnegative edge lengths $\ell \in \mathbb{R}_{\geq 0}^E$. For simplicity, we write $\ell_{v,w}$ instead of $\ell_{\{v,w\}}$ for $\{v,w\} \in E$. We consider the shortest path networks with respect to s_1 and s_2 . To that end, let $d_i: V \rightarrow \mathbb{R}_{\geq 0}$ denote the distance function induced by the lengths ℓ with respect to s_i for $i \in [2]$. The corresponding **shortest paths network** is given by the edge set

$$E_i := \{\{v, w\} \in E \mid |d_i(v) - d_i(w)| = \ell_{v,w}\}.$$

See [Figure 3.6a](#) for an example of the sets E_i .

For $i \in [2]$, the distance function d_i induces a natural orientation on all edges in E_i with strictly positive length. We orient all edges in E_i away from the source s_i ; that is, we replace an edge $\{v, w\} \in E_i$ by the arc (v, w) assuming $d_i(v) < d_i(w)$. The orientations induced by d_1 and d_2 on the edges in $E_1 \cap E_2$, however, may differ. Replacing a single edge with two opposite arcs interferes with the preservation of disjointness. Therefore, we substitute edges with ambiguous orientation by a (standard) gadget which preserves disjoint paths. It is depicted in [Figure 3.5](#). The gadget for an edge $\{v, w\} \in E$ contains a unique v - w path and a unique w - v path. These correspond to the two opposite orientations of $\{v, w\}$. Note that these two paths share a common arc in the gadget.




Figure 3.5: The gadget (right) for resolving the ambiguity in the orientation of an edge (left).

Hence, the gadget cannot be traversed by more than one arc/edge-disjoint path. Further, both paths between v and w in the gadget consist of exactly three arcs. Setting the length of all arcs the gadget to $\frac{1}{3}\ell_{v,w}$ preserves the distances between v and w . Thus, the distance functions d_i can be extended to the new vertices introduced with gadgets.

In the oriented graph, we distinguish between arcs the orientation of which is based on d_1 or d_2 . We denote these sets of arcs by A_1 and A_2 , respectively. For $i \in [2]$, the set A_i contains all arcs that result from orienting edges in E_i . Let $\{v, w\} \in E_i$ with $d_i(v) < d_i(w)$. For $\{v, w\} \in E_i$, the set A_i contains the arc (v, w) if and only if either $\{v, w\} \in E_1 \triangle E_2$ or the orientations induced by d_1 and d_2 agree. If $\{v, w\} \in E_1 \cap E_2$ and the orientations induced by d_1 and d_2 differ, A_1 and A_2 contain all arcs of the unique v - w path and the unique w - v path in the gadget that replaces $\{v, w\}$, respectively.

An orientation is induced by the distances for edges with strictly positive lengths only. Therefore, the set of edges with length zero $E_0 := \{e \in E \mid \ell_e = 0\}$ is left unoriented and dealt with separately.


This finishes the description of the graph transformation. We summarize it in the following definition.

Definition 3.5 (Partial orientation). Let $G = (V, E)$ be an undirected graph, $s \in V^2$ be a pair of sources and sinks, and $\ell \in \mathbb{R}_{\geq 0}^E$ be nonnegative edge lengths. The **partial orientation** of G with respect to s and ℓ is the graph $\vec{G} := (W, E_0 \cup A_1 \cup A_2)$ where W is the set of vertices V augmented with additional vertices introduced with gadgets, and E_0 , A_1 , and A_2 are as defined in the preceding paragraphs.

As described above, the distance functions of the original graph G can be extended to the augmented vertex set of \vec{G} . For $i \in [2]$ and $v \in W \setminus V$, we set $d_i(v)$ to the length of a shortest s_i - v path in \vec{G} , respectively.

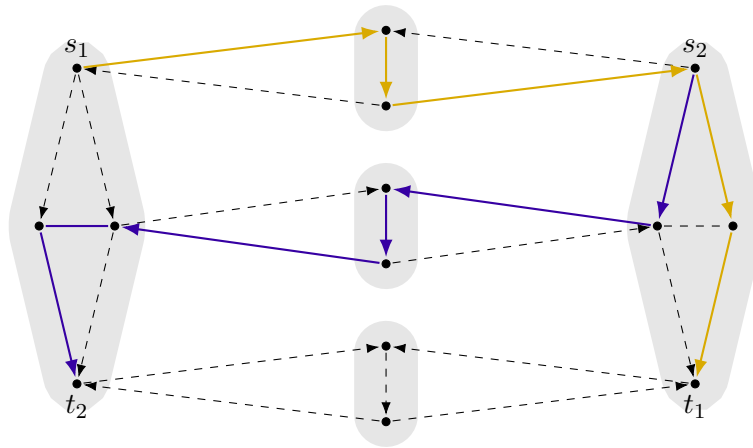
The partial orientation of the example from Figure 3.6a is depicted in Figure 3.6b. Its decomposition is illustrated in Figure 3.6c. As we discuss the existence of two arc/edge-disjoint paths restricted to different arc and edge sets in \vec{G} , we introduce a variant of the disjoint paths relation.

Definition 3.6 (Two disjoint paths relations). Let $G = (V, \mathcal{A})$ be a mixed graph, and $\mathcal{A}_1, \mathcal{A}_2 \subseteq \mathcal{A}$ be two subsets of arcs and edges. For $v, w \in V^2$, we write $v \overset{\mathcal{A}_1}{\rightleftarrows}_{\mathcal{A}_2} w$ if there exist a v_1 - w_1 path in \mathcal{A}_1 and a w_2 - v_2 path in \mathcal{A}_2 which are arc/edge-disjoint. If \mathcal{A}_1 and \mathcal{A}_2 agree, we write $\overset{\mathcal{A}_1}{\rightleftarrows}$ instead of $\overset{\mathcal{A}_1}{\rightleftarrows}_{\mathcal{A}_2}$.



(a) An undirected graph with annotated edge lengths. The shortest path networks E_1 and E_2 agree (solid edges).

(b) The graph's partial orientation consisting of the arcs A_1 (yellow), the arcs A_2 (violet), and the edges E_0 (black).



$$(d_1 - d_2)(\bullet) = -1 \quad (d_1 - d_2)(\bullet) = 0 \quad (d_1 - d_2)(\bullet) = 1$$

(c) The weakly connected components of $(W, E_0 \cup (A_1 \cap A_2))$ ordered nondecreasingly with respect to $d_1 - d_2$. The s_1 - t_1 path (yellow) traverses the components in that order. The s_2 - t_2 path (violet) visits them in the reverse order.

Figure 3.6: Exemplary construction of a partial orientation.

With the help of this notation, we can express the equivalence of the original and the transformed disjoint paths problems. The partial orientation does not only preserve shortest paths but isolates them. Edges outside the shortest path networks are not included. Note that the direction of the paths between the components of the related pairs is opposite.

Lemma 3.7 (Paths in the partial orientation). *Let $G = (V, E)$ be an undirected graph, $s \in V^2$ be a pair of sources, and $\ell \in \mathbb{R}_{\geq 0}^E$ be nonnegative edge lengths. Further, let $\vec{G} = (W, E_0 \cup A_1 \cup A_2)$ denote the partial orientation of G with respect to s and ℓ . Then for every $t \in V^2$, we have $s \xrightarrow{\ell}_E t$ if and only if $\binom{s_1}{t_2} \xleftrightarrow{E_0 \cup A_1}_{E_0 \cup A_2} \binom{t_1}{s_2}$.*

Proof. We transform edge-disjoint paths in G to the corresponding arc/edge-disjoint paths in \vec{G} by orienting edges, and vice versa.

Orienting paths. Assume there exist two edge-disjoint paths P_1 and P_2 such that P_i is a shortest s_i - t_i path with respect to ℓ in G for $i \in [2]$. By definition, $P_i \subseteq E_i$ holds for both $i \in [2]$. Every edge with nonzero length in P_i is replaced by the respective oriented arc or path in the respective gadget. We call the resulting path \vec{P}_i . Observe that \vec{P}_i is indeed a path in $E_0 \cup A_i$. This is due to P_i being a shortest path in E_i . Hence, the orientation of A_i agrees with the direction of P_i and \vec{P}_i . The two paths \vec{P}_1 and \vec{P}_2 are arc/edge-disjoint as different edges are replaced by disjoint (sets of) arcs.

Disorienting paths. Assume there are arc/edge-disjoint paths \vec{P}_1 and \vec{P}_2 such that \vec{P}_i is an s_i - t_i path in $E_0 \cup A_i$ for both $i \in [2]$. For every gadget that \vec{P}_i passes through, we replace the subpath of \vec{P}_i within the gadget by the corresponding edge in E_i . All arcs outside of gadgets are substituted by their respective undirected edges in E_i . We call the resulting path P_i . It is a shortest path in G , as \vec{P}_i uses only arcs and edges from $E_0 \cup A_i$. Further, the two obtained paths P_1 and P_2 must be edge-disjoint because \vec{P}_1 and \vec{P}_2 are arc/edge-disjoint. ■

3.4.2 Disjoint Paths in the Partial Orientation

The partial orientation uses orientation of the edges in order to enforce shortest paths as highlighted by Lemma 3.7. The paths of the two source-sink pairs are restricted to the two different sets of arcs and edges $E_0 \cup A_1$ and $E_0 \cup A_2$. In order to obtain disjointness of the two paths, we need to pay attention to the intersection of these sets, that is, $E_0 \cup (A_1 \cap A_2)$. The following lemma sheds light on the special structure of this intersection which we can exploit algorithmically.

Lemma 3.8 (Structure of the partial orientation). *Let $G = (V, E)$ be an undirected graph, $s \in V^2$ be a pair of sources, and $\ell \in \mathbb{R}_{\geq 0}^E$ be nonnegative edge lengths. Further, let $\vec{G} = (W, E_0 \cup A_1 \cup A_2)$ be the partial orientation of G with respect to s and ℓ . If $W = \bigcup_{j=1}^h W_j$ is the partition of the augmented vertex set W into the vertex sets of the h weakly connected components of the subgraph $(W, E_0 \cup (A_1 \cap A_2))$, then*

- (i) $\vec{G}[W_j]$ is weakly acyclic for all $j \in [h]$,
- (ii) $\vec{G}[W_j]$ contains no arcs from $A_1 \triangle A_2$ for all $j \in [h]$,
- (iii) sorting the connected components $W_j, j \in [h]$ in \vec{G} nondecreasingly with respect to $d_1 - d_2$ simultaneously yields a topological ordering of $(W, E_0 \cup A_1) / \{W_1, \dots, W_h\}$ and a reverse topological ordering of $(W, E_0 \cup A_2) / \{W_1, \dots, W_h\}$.

Algorithm 2: Dynamic program for the two disjoint shortest paths problem with nonnegative edge lengths

Input: undirected graph $G = (V, E)$, pairs of sources and sinks $s, t \in V^2$,

nonnegative edge lengths $\ell \in \mathbb{R}_{\geq 0}^E$

Output: $s \xrightarrow{\ell}_E t$

- 1 Construct $\vec{G} = (W, E_0 \cup A_1 \cup A_2)$ for G with respect to s and ℓ ;
 - 2 Find weakly connected components W_1, \dots, W_h of $(W, E_0 \cup (A_1 \cap A_2))$ sorted nondecreasingly with respect to $d_1 - d_2$;
 - 3 **for** $j = 1, \dots, h$ **do** compute $\xleftrightarrow{\vec{G}[W_j]}$ and $\xleftrightarrow{\delta_{A_1}^-(W_j)} \delta_{A_2}^+(W_j)$;
 - 4 Initialize $\xleftrightarrow{\quad}$ to the identity relation $\{(v, v) \mid v \in W^2\}$;
 - 5 **for** $j = 1, \dots, h$ **do** update $\xleftrightarrow{\quad}$ to $\xleftrightarrow{\vec{G}[W_j]} \circ \xleftrightarrow{\delta_{A_1}^-(W_j)} \delta_{A_2}^+(W_j) \circ \xleftrightarrow{\quad}$;
 - 6 **return** $s \xleftrightarrow{\quad} t$;
-

Proof. We deduce the statements from the monotonicity of d_1 and d_2 along arcs and edges. Consider the function $d_1 - d_2$ on the vertex set of \vec{G} . For $(v, w) \in A_1 \setminus A_2$, the definitions of the arc sets give $d_1(w) = d_1(v) + \ell_{v,w}$ and $d_2(w) < d_2(v) + \ell_{v,w}$. Subtracting these (in)equalities yields that $d_1 - d_2$ is strictly increasing along arcs in $A_1 \setminus A_2$. By symmetry, $d_1 - d_2$ strictly decreases along arcs in $A_2 \setminus A_1$. It can be shown in a similar manner, that $d_1 - d_2$ is constant along edges in E_0 and arcs in $A_1 \cap A_2$. As an immediate consequence, it is constant on W_j for every $j \in [h]$ and the statement (ii) is true. Together these monotonicities of $d_1 - d_2$ imply (iii). The distance functions d_1 and d_2 are strictly increasing along arcs in $A_1 \cap A_2$, while they are constant along edges in E_0 . Thus, there cannot be a cycle within $\vec{G}[W_j]$ that contains an arc from $A_1 \cap A_2$. The characterization following Definition 3.3 then implies the statement (i). ■

This structural result allows to apply Theorem 3.2 to the two disjoint paths relation on \vec{G} . Strictly speaking, the relation $\xleftrightarrow{E_0 \cup A_1}^{E_0 \cup A_2}$ is more specific than the disjoint paths relation covered by the theorem. Its statement still extends as the partition and concatenation of paths preserves them being part of $E_0 \cup A_1$ or $E_0 \cup A_2$. We obtain that $\xleftrightarrow{E_0 \cup A_1}^{E_0 \cup A_2}$ can be decomposed into

$$\xleftrightarrow{\vec{G}[W_h]} \circ \xleftrightarrow{\delta_{A_1}^-(W_h)} \delta_{A_2}^+(W_h) \circ \xleftrightarrow{\vec{G}[W_{h-1}]} \circ \dots \circ \xleftrightarrow{\delta_{A_1}^-(W_3)} \delta_{A_2}^+(W_3) \circ \xleftrightarrow{\vec{G}[W_2]} \circ \xleftrightarrow{\delta_{A_1}^-(W_2)} \delta_{A_2}^+(W_2) \circ \xleftrightarrow{\vec{G}[W_1]}.$$

Algorithm 2 therefore works in a very similar way to Algorithm 1. The arcs and edges that both (oriented) shortest path networks have in common are treated separately from




Figure 3.7: In the j -th iteration of Line 5 in [Algorithm 2](#), the relation \rightleftharpoons is extended by concatenating previously computed paths with arcs from the disjoint sets $\delta_{A_1}^-(V_j)$ and $\delta_{A_2}^+(V_j)$, and arc/edge-disjoint paths in $G[V_j]$. The path for the first component is extended by appending to its end. The path for the second component is extended by prepending to its beginning.

their individual parts. [Lemma 3.8](#) shows that the common parts are weakly acyclic mixed graphs which are connected by a directed acyclic structure when restricting to A_1 or A_2 . Therefore, [Algorithm 1](#) is applied to compute the disjoint paths relation on the common parts. Again, the disjoint paths relation between them have a very simple structure. Dynamic programming is applied to compose the relations on all parts to compute $\rightleftharpoons_{E_0 \cup A_2}^{E_0 \cup A_1}$ on successively larger parts of \vec{G} ; see [Figure 3.7](#). Note that we do not fully recover $\rightleftharpoons_E^\ell$ as the partial orientation \vec{G} is constructed with respect to s .

Theorem 3.9 (Correctness and complexity). *Given an undirected graph $G = (V, E)$ with nonnegative edge lengths $\ell \in \mathbb{R}_{\geq 0}^E$ and $s, t \in V^2$, [Algorithm 2](#) decides $s \xrightarrow{\ell}_E t$ in polynomial time.*

Proof. Let $W = \bigcup_{j=1}^h W_j$ be the partition of W into the vertex sets of the h weakly connected components of the subgraph $(W, E_0 \cup (A_1 \cap A_2))$ as computed by the algorithm. [Lemma 3.8 \(iii\)](#) shows that the W_j 's are simultaneously sorted in topological order with respect to $(W, E_0 \cup A_1) / \{W_1, \dots, W_h\}$ and in reverse topological order with respect to $(W, E_0 \cup A_2) / \{W_1, \dots, W_h\}$.

Correctness. For $i \in [2]$ and $j \in \{0, \dots, h\}$, set $\mathcal{A}_i^{(j)}$ to be the arcs of A_i and edges of E_0 in the induced subgraph $\vec{G}[\bigcup_{l=1}^j W_l]$. In particular, we have $\mathcal{A}_i^0 = \emptyset$. As already discussed [Theorem 3.2](#) can be extended in a straightforward manner to apply to $\rightleftharpoons_{E_0 \cup A_2}^{E_0 \cup A_1}$.

It follows by induction that \rightleftharpoons equals $\rightleftharpoons_{\mathcal{A}_2^{(j)}}^{\mathcal{A}_1^{(j)}}$ after the j -th iteration of Line 5. Therefore, it is $\rightleftharpoons_{\vec{G}}$ in Line 6. [Lemma 3.8](#) now implies correctness of [Algorithm 2](#).

Complexity. As for the running time, finding the weakly connected components and sorting them nondecreasingly with respect to $d_1 - d_2$ can be done in polynomial time. Computing the relations $\rightleftharpoons_{G[W_j]}$ also can be done efficiently as we have shown in [Theorem 3.4](#). Finally, the binary relations on V^2 have at most $|V|^4$ elements. Therefore, composing a polynomial number of them takes only polynomial time. In total, the running time of the algorithm is polynomial in the input size. \blacksquare

Similar to [Algorithm 1](#), also [Algorithm 2](#) can be adapted to decide the existence of two vertex-disjoint shortest paths in undirected graphs. The changes are the same. The construction of the partial orientation can be even simplified. The gadget in [Figure 3.5](#) is not necessary. It suffices to introduce two opposite arcs instead as only one of them can be used by two vertex-disjoint paths.

3.5 Closing Remarks

In this chapter, we facilitate an acyclic decomposition of graphs in order to develop efficient algorithms for disjoint paths problems. Applying this decomposition twice allows us to solve the two disjoint shortest paths problem in polynomial time. Interestingly, the known results on the undirected disjoint shortest paths problem do not show a separation of the complexities for $k = 2$ and fixed $k \geq 3$. The techniques used in [Section 3.4](#), however, do not seem to extend beyond $k = 2$. They heavily rely on the acyclicity of $(W, E_0 \cup A_1) / \{W_1, \dots, W_h\}$ and $(W, E_0 \cup A_2) / \{W_1, \dots, W_h\}$ as established by [Lemma 3.8](#). It is not clear what a corresponding decomposition would look like for $k = 3$.

The most apparent unsolved questions on the complexity of the disjoint shortest paths problem are represented by the gaps in [Table 3.1](#). In terms of the treatment of zero-length edges, a first step might be combining our techniques with those from Bentert et al. (2021) and Lochet (2021). This might allow to tackle instances of binary edge lengths. The case $\ell \geq 0$ seems to necessitate the understanding of arbitrary positive edge lengths first for $k \geq 3$. Another interesting open problem is the shortest two disjoint paths problem beyond the case $\ell \equiv 1$.

Chapter 4

Dynamic Equilibria under the Fluid Queuing Network

In this chapter, we consider **dynamic equilibria** for flows over time under the **fluid queuing model**. In this model, queues on the links of a network take care of flow propagation. Flow enters the network at a single source and leaves at a single sink. In a dynamic equilibrium, every infinitesimally small flow particle reaches the sink as early as possible given the pattern of the rest of the flow. While this model has been examined for many decades, progress has been relatively recent. In particular, the derivatives of dynamic equilibria have been characterized as **thin flows with resetting**, which allowed for more structural results.

Our two main results are based on the formulation of thin flows with resetting as **linear complementarity problem** and its analysis. We present a constructive proof of existence for dynamic equilibria if the inflow rate is **right-monotone**. The complexity of computing thin flows with resetting, which occurs as a subproblem in this method, is still open. We settle it for the class of **series-parallel** networks by giving a recursive algorithm that solves the problem for all flow values simultaneously in polynomial time.

Authorship. This work has been previously published by the author of this thesis (2022). The presentation in this chapter coincides largely with the article.

Outline. This chapter is organized as follows. [Section 4.2](#) introduces the fluid queuing model and the notion of a dynamic equilibrium in a more formal manner. The most important known results for our purposes are stated. This is continued in [Section 4.3](#) in which the concept of a normalized thin flows with resetting is defined, and known characterizations are given. In [Section 4.3](#), we further examine normalized thin flows with resetting parametrized by the flow value, which is needed in the subsequent two sections. [Section 4.4](#) generalizes the known constructive proof of the existence of dynamic equilibria to be able to cope with right-monotone, locally integrable inflow rates. In [Section 4.5](#), we consider the computation of normalized thin flows with resetting. We prove that this computation can be done efficiently for two-terminal series-parallel networks. Finally, [Section 4.6](#) finishes with some concluding remarks.

4.1 Introduction

Systems with large temporal fluctuation, as it emerges for example in road traffic, cannot be adequately captured by models based on static flows. This shortcoming is already addressed by Ford and Fulkerson (1962), who look at flow that travels through a network at a (finite) speed. In their basic model for **flows over time**, the network is defined on a directed graph. Each arc represents a link with some constant delay and capacity. Many optimization problems in this model are well-understood.

The **fluid queuing model**, as considered already by Vickrey (1969) in the context of transportation investments, builds upon this model and allows to study flow which exhibits selfish behavior. Each link in the network is equipped with a fluid queue at its entrance. If the capacity constraint for a link is violated, the excess flow is collected in its queue. Consequently, waiting in the queue imposes additional delay for flow to traverse a link. For a given inflow rate at the source, each infinitesimally small flow particle is considered a player in a **routing game over time** who tries to reach the sink in the shortest possible time. Under this assumption, a **dynamic equilibrium** describes a state of the system in which no player has an incentive to deviate. Each player is assumed to have full information. Hence, they anticipate the behavior of other flow particles and the state of the fluid queues at the respective moment they would reach the queues. As the queues respect the first in, first out principle, there cannot be overtaking of particles in a dynamic equilibrium. Therefore, the shortest path for a particle entering at the source is determined by all the flow which has entered before.

4.1.1 Related Literature

We give an overview of the literature on the topic. Before summarizing the existing findings on the fluid queuing model, we present some major results on optimization problems dealing with flows over time.

Flow over time. Many optimization problems in the context of flow over time are well understood. Computing the maximum amount of flow that can be sent from a source to a sink within a given time horizon is efficiently solved by Ford and Fulkerson (1958). The authors show that the **maximum s - t flow over time** problem always has an optimal solution with a simple structure. It is based on a static path flow and is called temporally repeated flow.

Switching the roles of the flow value and time horizon yields another problem which is the **quickest s - t flow** problem. It asks for the smallest time horizon in which a given amount of flow can be sent from a source to a sink. An optimal temporally repeated flow to this problem can be found by a parametric search and iterative maximum s - t flow over time computations. Burkard, Dlaska, and Klinz (1993) obtain a strongly polynomial algorithm by applying the discrete Newton method of Megiddo (1978, 1979). The fastest currently known strongly polynomial algorithm is due to Saho and Shigeno (2017). Allowing more sources and sinks (of a single commodity) results in the **quickest transshipment problem**. Hoppe and Tardos (1995, 2000) show that there are still optimal

(generalized) temporally repeated flows for this problem when the path decomposition of the underlying static flow allows paths in the residual network. Their algorithm extensively uses parametric submodular minimization, which results in a high running time. This is significantly improved by Schlöter, Skutella, and Van Tran (2021).

An s - t flow over time which is a maximum s - t flow over time for all time horizons (or equivalently a quickest s - t flow for all flow values) is called an **earliest arrival s - t flow**. Such flows over time are particularly interesting in the context of evacuation scenarios. Gale (1959) shows their existence. Exact algorithms are given by Minieka (1973) and Wilkinson (1971). The natural encoding of the output, however, can have exponential size as illustrated by an example of Zadeh (1973). Therefore, it is unlikely that exact polynomial-time algorithms exist. For this reason, the fully polynomial approximation scheme by Hoppe and Tardos (1994) seems to be the most efficient possible. An exact algorithm for the **earliest transshipment problem** in PSPACE is found by Schlöter and Skutella (2017).

Computing a **minimum-cost flow over time** on the other hand is NP-hard even in series-parallel networks as shown by Klinz and Woeginger (1995, 2004). Based on the idea of **condensed time expanded networks**, Fleischer and Skutella (2007) devise a fully polynomial approximation scheme.

Many of the mentioned results have been first obtained for a variant of the model with discrete time steps. The corresponding results for the continuous time model are due to Anderson and Philpott (1994) and Fleischer and Tardos (1998). For more details on flows over time, see the introductory article by Skutella (2009) and the survey by Kotnyek (2003).

Dynamic traffic assignment. Traffic planners tend to base their decisions on complex simulations. The theory regarding the existence and computation of dynamic equilibria, however, has not caught up on the complexity of the models in use (see the report by Cominetti, Harks, Osorio, and Peis (2018)). The broad class of computational problems related to traffic models is referred to as **traffic assignment problems**, as introduced by Merchant and Nemhauser (1978a,b). In particular, it contains the problem at hand. An overview of the topic is available by Peeta and Ziliaskopoulos (2001).

Dynamic equilibria under the fluid queuing model. The concept of dynamic equilibria dates back to Vickrey (1969) and Yagar (1971). Existence of dynamic equilibria for a class of models related to the one described above was shown by Zhu and Marcotte (2000). The work by Meunier and Wagner (2010) concerns a general framework for dynamic congestion games. Their results regarding the existence of equilibria are the first to apply to the model as defined above. The involved methods, however, are nonconstructive. The basis for a constructive proof of existence of dynamic equilibria and further insights into their nature is provided by Koch and Skutella (2009, 2011). These authors establish a specific structure of the derivatives of dynamic equilibria which they call **thin flows with resetting**. Based on this, they suggest a method to compute a dynamic equilibrium for constant inflow rates by integrating over thin flows with resetting. The integration

is done in phases during which the derivative stays constant. One such step is called **α -extension**.

Cominetti, Correa, and Larré (2011, 2015) refine the notion of thin flows with resetting to **normalized thin flows with resetting** and prove their existence and the uniqueness of the associated labels. This turns the α -extension into a constructive proof of existence of dynamic equilibria for piecewise constant inflow rate. The authors augment this by also showing existence for inflow rates in $L^p(0, T)$, where $1 < p < \infty$, $T \geq 0$, via variational inequalities (also for multiple origin-destination pairs). Sering and Skutella (2018) obtain an algorithmic approach for a multiple-source multiple-sink setting, in which the inflow rates at the sources are constant and the inflow of every source is routed to the sinks proportionally with respect to a global pattern. Pham and Sering (2020) extend the results to the setting of time-dependent capacities and transit times.

Behavior of dynamic equilibria. The model allows notions of a **price of anarchy** with respect to various objectives for the social optimum. First results are given by Koch and Skutella (2009, 2011). They show that the price of anarchy with respect to a maximum flow over time is unbounded.

Bhaskar, Fleischer, and Anshelevich (2011, 2015) analyze the price of anarchy with respect to a quickest flow over time in a Stackelberg game. By reducing the capacities of the links, a leader can ensure that a fixed amount of flow in a dynamic equilibrium reaches the sink within a factor of $\frac{e}{e-1}$ compared to the time it would take in a social optimum of the original network. Further progress on this price of anarchy is made by Correa, Cristi, and Oosterwijk (2019), who prove a bound of $\frac{e}{e-1}$ under weaker assumptions, which get dispensable if a certain **monotonicity conjecture** holds.

For static routing games, Braess (1968) shows that the introduction of additional links in a network can decrease the quality of an equilibrium. This phenomenon is since known as **Braess's paradox**. Macko, Larson, and Steskal (2010, 2013) show that the model at hand exhibits this behavior as well even in series-parallel networks.

Cominetti, Correa, and Olver (2017) consider the long-term behavior of dynamic equilibria. They show that for constant inflow rates the queues stay constant after a finite amount of time (given a necessary condition). At the same time, they give a small series-parallel network capturing distinctive phenomena of dynamic equilibria. In particular, the flow in a dynamic equilibrium over some cut can exceed the constant inflow rate (within a bounded time interval).

Variants of the model. Methods similar to the ones applied to the model at hand were successfully used on related models. Sering and Vargas Koch (2019) investigate a variant of the fluid queuing model in which the total amount of flow on a link at any point in time is bounded. If links fill up, congestion is propagated backwards across vertices to their incoming links, which is called **spillback**. Israel and Sering (2020) examine the impact of this model augmentation on the price of anarchy.

Graf and Harks (2019) and Graf, Harks, and Sering (2020), on the other hand, achieve existence results for a different concept of so-called **instantaneous dynamic equilibria**

in the fluid queuing model. It captures the scenario in which players do not have full information.

4.1.2 Our Contribution

We linearize a known (nonlinear) complementarity problem by Cominetti, Correa, and Larré (2011) and use it to examine normalized thin flows with resetting and their dependency on the flow value (which is determined by the inflow rate). This enables us to generalize the α -extension to a larger class of inflow rates, namely right-monotone, locally integrable functions. Here, a function is **right-monotone** if each point is the lower end point of a closed interval on which the function is monotone. We also use our new insights into the properties of normalized thin flows with resetting to tackle their computation. This subproblem in the α -extension is known to be in PPAD, and it remains unclear whether it lies in P. We settle the computational complexity for the class of **series-parallel** networks by giving a polynomial-time algorithm that solves the problem for all flow values simultaneously.

4.2 The Fluid Queuing Model and Dynamic Equilibria

Throughout this chapter, we consider a network which is given by a directed graph $G = (V, A)$, positive arc capacities $\nu \in \mathbb{R}_{>0}^A$, and nonnegative delays $\tau \in \mathbb{R}_{\geq 0}^A$. We assume that there is no (directed) cycle $C \subseteq A$ in G with zero total transit time, that is, $\tau(C) = 0$. An exemplary network is depicted in Figure 4.2.

Let $s \in V$ be the source and $t \in V$ be the sink of a single commodity. We assume that there is a (directed) s - v path in G for every vertex $v \in V$. This is without loss of generality, as flow emanating from s can only reach these vertices. Further, the rate at which flow enters the network at the source s at any point in time is described by a nonnegative function $\nu_0 \in L_{\text{loc}}^1(\mathbb{R})$ that vanishes almost everywhere on $\mathbb{R}_{<0}$. Let $N_0: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ denote the cumulative inflow, that is, $N_0(\vartheta) = \int_0^\vartheta \nu_0(\theta) d\theta$ for every $\vartheta \in \mathbb{R}$. Then, N_0 is locally absolutely continuous, that is, it is absolutely continuous on every bounded interval.

A **flow over time** in the network is defined by two functions $f^+, f^-: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^A$ which vanish almost everywhere on $\mathbb{R}_{<0}$. For any $a \in A$, we denote by f_a^+ the function that maps $\vartheta \in \mathbb{R}$ to the component a of $f^+(\vartheta)$, that is, $f_a^+(\vartheta) = (f^+(\vartheta))_a$. An analogous notation is used for f^- and other vector-valued functions. For $a \in A$, the functions f_a^+ and f_a^- have to be locally integrable and represent the inflow and outflow rates of arcs, respectively, in dependency on the time. Define the cumulative flow functions $F^+, F^-: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^A$ by

$$F_a^+(\vartheta) := \int_0^\vartheta f_a^+(\theta) d\theta \quad \text{and} \quad F_a^-(\vartheta) := \int_0^\vartheta f_a^-(\theta) d\theta \quad \text{for all } a \in A \text{ and } \vartheta \in \mathbb{R}.$$

In the following, we describe the dynamics of the fluid queuing model, that is, the conditions a flow over time has to satisfy in order to obey the model. Each link represented




Figure 4.1: A link (left) represented by an arc $a = (v, w) \in A$ with transit time τ_a , capacity ν_a , inflow rate f_a^+ , outflow rate f_a^- , queue length z_a , and queuing delay q_a . We use a simplified representation (right).

by an arc $a \in A$ has a **fluid queue** at its entrance. The inflow and the outflow rate of a are related by the dynamics of that queue; see Figure 4.1. Flow that enters the link first has to pass through the queue before it may traverse it. The capacity of the arc limits the rate at which flow can leave the queue. After leaving, it takes the flow exactly τ_a units of time to traverse the link. We assume that queues initially are empty and **operate at capacity**, that is, as much flow as the capacity permits leaves the queue. Let $z_a: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ be the function that describes the cumulative flow which resides in the queue of a in dependency on the time. Therefore, z_a is the unique locally absolutely continuous solution (see, e.g., Filippov 1988, p. 106) to

$$z_a \equiv 0 \text{ on } \mathbb{R}_{\leq 0} \quad \text{and} \quad \frac{dz_a}{d\vartheta}(\vartheta) = \begin{cases} [f_a^+(\vartheta) - \nu_a]_+ & \text{if } z_a(\vartheta) \leq 0 \\ f_a^+(\vartheta) - \nu_a & \text{if } z_a(\vartheta) > 0 \end{cases} \quad \text{for a.e. } \vartheta \geq 0. \quad (\text{QD})$$

As flow takes exactly τ_a units of time after leaving the queue to traverse a , we require for all times $\vartheta \in \mathbb{R}$ that $z_a(\vartheta) = F_a^+(\vartheta) - F_a^-(\vartheta + \tau_a)$ holds. Together with (QD) this determines the outflow rate of $a \in A$ as

$$f_a^-(\vartheta + \tau_a) = f_a^+(\vartheta) - \frac{dz_a}{d\vartheta}(\vartheta) = \begin{cases} \min \{ f_a^+(\vartheta), \nu_a \} & \text{if } z_a(\vartheta) = 0 \\ \nu_a & \text{if } z_a(\vartheta) > 0 \end{cases} \quad \text{for a.e. } \vartheta \in \mathbb{R}.$$

Finally, we impose **strict flow conservation** at any vertex $v \in V \setminus \{t\}$ but the sink, that is,

$$\sum_{a \in \delta^+(v)} f_a^+(\vartheta) - \sum_{a \in \delta^-(v)} f_a^-(\vartheta) = \begin{cases} \nu_0(\vartheta) & \text{if } v = s \\ 0 & \text{if } v \in V \setminus \{s, t\} \end{cases} \quad \text{for a.e. } \vartheta \in \mathbb{R}.$$

Any flow over time which obeys the queuing dynamics and strict flow conservation is called **feasible** in the following.

The above defines the model from a cumulative point of view. While it describes feasible behavior of flow in its entirety, it does not keep track of single flow particles. In the routing game that we are about to consider, each infinitesimally small flow particle is considered a player and, therefore, must be trackable. In order to allow that, the **first in, first out principle** is imposed on the queues. By doing so, the **queuing delay** on

4.2 The Fluid Queuing Model and Dynamic Equilibria




Figure 4.2: An exemplary network with transit times τ and capacities ν . The line widths are proportional to the capacities. A constant inflow of rate $\nu_0 \equiv 1$ arrives at the source s and travels towards the sink t .

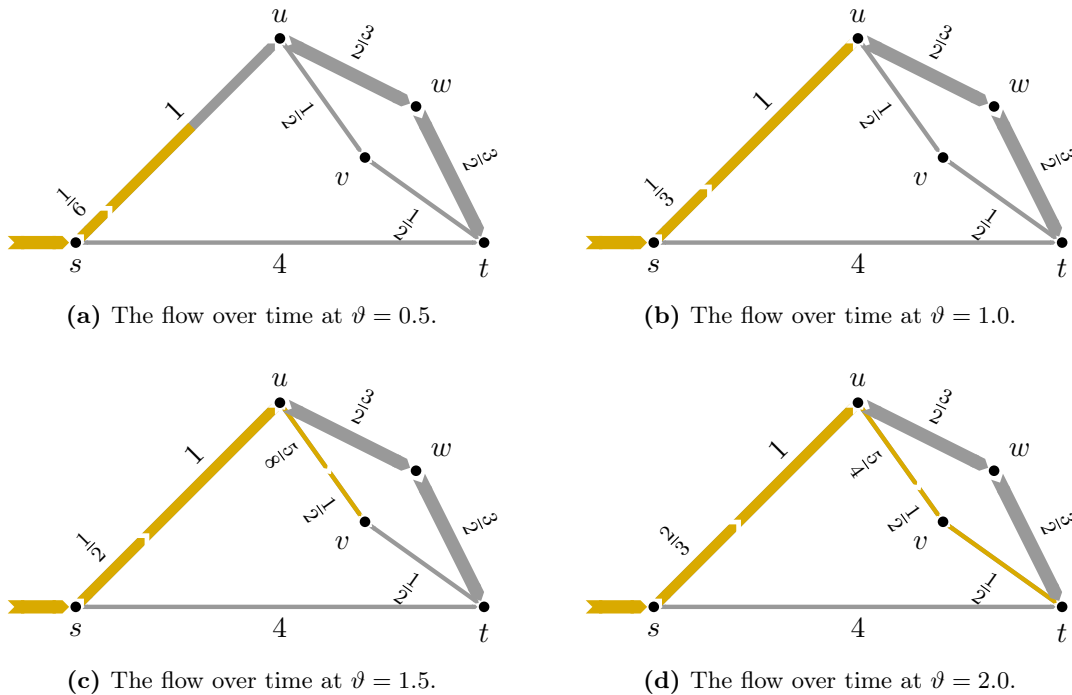


Figure 4.3: A flow over time along the path s - u - v - t through the network depicted in Figure 4.2. Every arc a is annotated with its respective queuing delay $q_a(\vartheta)$ (if it is nonzero) and its transit time τ_a . The relative contributions of $q_a(\vartheta)$ and τ_a to the total transit time $q_a(\vartheta) + \tau_a$ is indicated by a small gap in the link.

an arc $a \in A$ which is experienced by flow entering it at time $\vartheta \in \mathbb{R}$ is defined as

$$q_a(\vartheta) := \min\{q \geq 0 \mid F_a^-(\vartheta + \tau_a + q) \geq F_a^+(\vartheta)\}.$$

For queues operating at capacity, this evaluates to $q_a(\vartheta) = z_a(\vartheta)/\nu_a$. Thus, a particle which enters an arc $a \in A$ at time ϑ leaves it exactly at time $T_a(\vartheta) := \vartheta + q_a(\vartheta) + \tau_a$. Based on strict flow conservation at the vertices, flow that enters a path $P = \{a_1, \dots, a_{|P|}\}$ at time ϑ reaches its end exactly at time $\ell^P(\vartheta) := T_{a_{|P|}} \circ T_{a_{|P|-1}} \circ \dots \circ T_{a_1}(\vartheta)$. Let \mathcal{P}_{sv} denote the set of all s - v paths in G . The **earliest time** that a particle, which leaves the source at time $\vartheta \in \mathbb{R}$, can reach v is $\ell_v(\vartheta) := \min_{P \in \mathcal{P}_{sv}} \ell^P(\vartheta)$. Every path attaining the minimum is called a **dynamic shortest s - v path** relative to time ϑ .

Switching from this path-based definition to an arc-based view, $\ell: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^V$ is determined by the **Bellman equations**

$$\ell_s(\vartheta) = \vartheta \quad \text{and} \quad \ell_w(\vartheta) = \min_{\substack{a \in \delta^-(w) \\ a=(v,w)}} T_a(\ell_v(\vartheta)) \quad \text{for all } w \in V \setminus \{s\} \text{ and } \vartheta \in \mathbb{R}. \quad (\text{BE})$$

The arcs for which the minima in (BE) are attained are called **active** relative to time ϑ . An s - v path in G that uses only active arcs relative to ϑ is also a dynamic shortest path relative to ϑ . Another type of arcs plays an important role in the model. An arc $(v, w) \in A$ with nonzero queue length at time $\ell_v(\vartheta)$ is called **resetting** relative to time ϑ . Note that q, T , and ℓ depend on the flow over time (f^+, f^-) . Further, f^+ can be recovered from f^+ and vice versa.

A feasible flow over time through the network from Figure 4.2 is illustrated in Figure 4.3. All flow travels along the path s - u - v - t . As the capacity of (s, u) is strictly below the inflow rate, a queue keeps growing with constant rate. The arc (u, v) gives another bottleneck, which is why it also grows a queue.

Interpreting each infinitesimally small flow particle in an s - t flow over time as a player who wants to minimize her travel time from s to t defines a routing game over time. A **dynamic equilibrium** is a flow over time in which no player can strictly decrease her travel time by deviating on her own. A flow particle that starts at s at time $\vartheta \in \mathbb{R}$ only uses active arcs relative to ϑ and, therefore, reaches every vertex v on its path from s to t as early as possible at time $\ell_v(\vartheta)$. That way, no particle starting at s after time ϑ can overtake and delay it. As players travel on dynamic shortest paths only, the dynamic equilibrium follows the first principle of Wardrop (1952) and, thus, corresponds to a user equilibrium of flow over time under the fluid queuing model.

Definition 4.1 (Dynamic equilibrium). *A feasible flow over time (f^+, f^-) is a **dynamic equilibrium** if for each arc $a = (v, w) \in A$ the inflow rate f_a^+ vanishes almost everywhere on the set $\{\ell_v(\vartheta) \mid a \text{ inactive relative to } \vartheta \in \mathbb{R}\}$.*

The flow over time in Figure 4.3 is not a dynamic equilibrium. For small $\vartheta \geq 0$, both arcs on the path $P = \{(s, u), (u, v), (v, t)\}$ are active. The queues on (s, u) and (u, v) grow arbitrarily large as ϑ tends to infinity. Thus, $\ell^P(\vartheta) - \ell_s(\vartheta) = +\infty$ for $\vartheta \rightarrow +\infty$. Due to the empty arc (s, t) , however, the earliest arrival time at t is bounded by $\ell_t(\vartheta) - \ell_s(\vartheta) \leq 4$

4.2 The Fluid Queuing Model and Dynamic Equilibria

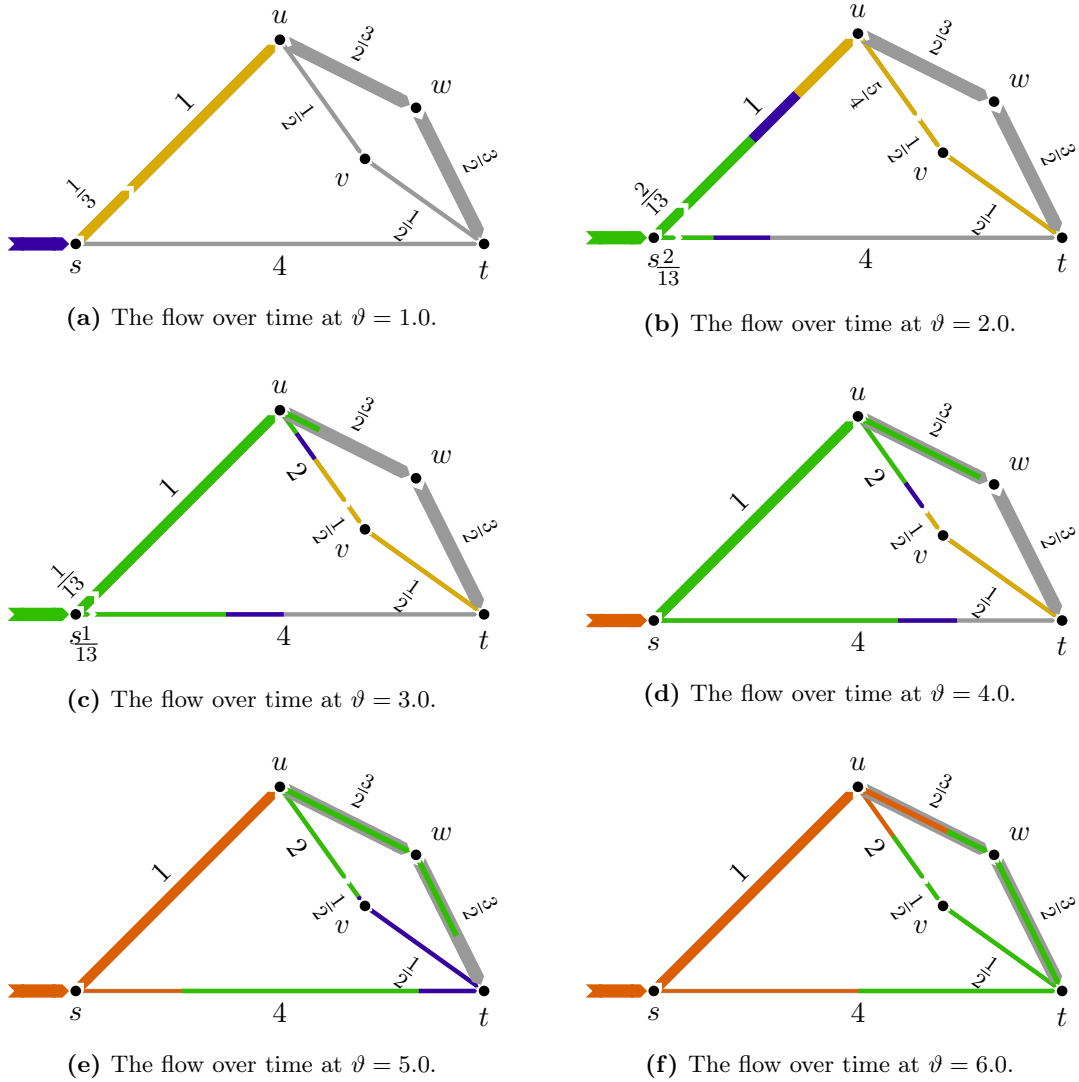


Figure 4.4: A dynamic equilibrium in the network depicted in Figure 4.2. The flow is colored with respect to its phases. Within each phase, the flow pattern is the same. Every arc a is annotated with its respective queuing delay $q_a(\vartheta)$ (if it is nonzero) and its transit time τ_a . The relative contribution of $q_a(\vartheta)$ to the total transit time $q_a(\vartheta) + \tau_a$ is indicated by a small white gap in the link.

for all $\vartheta \geq 0$. Hence, at some time (which is at $\vartheta = 1$) one of the arcs on P becomes inactive. Figure 4.4 shows a dynamic equilibrium in the same network. The coloring of the flow highlights different behavior and becomes clear in the further discussion. Flow leaving at the source in the time interval $\vartheta = [0, 1)$ (yellow) does take the same path as all flow in Figure 4.3. Flow particles that leave the source at times $\vartheta = [1, 1.4)$ (violet), however, travel along the arc (s, t) additional to the path P . After time $\vartheta = 1.4$ all arcs

of the network are active. During the phase $\vartheta \in [1.4, 4.0)$ (green) the queue on the arc (s, u) depletes. In the final phase (red) all queues are constant.

Definition 4.1 agrees with the one given by Cominetti, Correa, and Larré (2015). They show that dynamic equilibria can be characterized in terms of cumulative flows and the earliest times function as stated in the next lemma.

Lemma 4.2 (Characterization of dynamic equilibria). *A feasible flow over time (f^+, f^-) is a dynamic equilibrium if and only if $F_a^+(\ell_v(\vartheta)) = F_a^-(\ell_w(\vartheta))$ for all $a = (v, w) \in A$ and $\vartheta \in \mathbb{R}$.*

For a flow over time with earliest times function ℓ and a time $\vartheta \in \mathbb{R}$, we define the two sets of arcs

$$A'_\vartheta := \left\{ a = (v, w) \in A \mid \ell_w(\vartheta) \geq \ell_v(\vartheta) + \tau_a \right\} \quad \text{and}$$

$$A^*_\vartheta := \left\{ a = (v, w) \in A \mid \ell_w(\vartheta) > \ell_v(\vartheta) + \tau_a \right\}.$$

The following descriptions of these arc sets by Cominetti, Correa, and Larré (2015) give more intuition of their meaning. If ℓ is the earliest times function of a dynamic equilibrium, A'_ϑ and A^*_ϑ are exactly the sets of active arcs relative to ϑ and resetting arcs relative to ϑ , respectively, that is,

$$A'_\vartheta = \left\{ a = (v, w) \in A \mid \ell_w(\vartheta) = T_a(\ell_v(\vartheta)) \right\} \quad \text{and}$$

$$A^*_\vartheta = \left\{ a = (v, w) \in A \mid z_a(\ell_v(\vartheta)) > 0 \right\}.$$

A'_ϑ and A^*_ϑ are typically defined by these characterizations. We use the definition based on ℓ only as we will need these arc sets for flows over time which are not in equilibrium and, hence, for which the characterization does not hold.

A'_ϑ and A^*_ϑ are acyclic, as G does not contain cycles with total transit time zero. For dynamic equilibria, we additionally know that A'_ϑ contains an s - v path for every $v \in V$. This follows from the above characterization of A'_ϑ in a dynamic equilibrium and the Bellman equations (BE). This makes the triple $G'_\vartheta := (V, A'_\vartheta, A^*_\vartheta)$ a **shortest path graph with resetting** in the following sense.

Definition 4.3 (Shortest path graphs with resetting). *A shortest path graph with resetting is a triple (V, A', A^*) such that $A^* \subseteq A' \subseteq V \times V$ holds, A' is acyclic, and A' contains an s - v path for all $v \in V$.*

4.3 Normalized Thin Flows with Resetting

This section summarizes and extends some properties of the derivatives of dynamic equilibria. Assume that (f^+, f^-) is a dynamic equilibrium. We define $x: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^A$ by $x_a \equiv F_a^+ \circ \ell_v$ for all $a = (v, w) \in A$. From the strict flow conservation and Lemma 4.2,

it follows that for every $\vartheta \in \mathbb{R}$, the vector $x(\vartheta)$ is a static s - t flow in G as it satisfies

$$\sum_{a \in \delta^+(v)} x_a(\vartheta) - \sum_{a \in \delta^-(v)} x_a(\vartheta) = \sum_{a \in \delta^+(v)} F_a^+(\ell_v(\vartheta)) - \sum_{a \in \delta^-(v)} F_a^-(\ell_v(\vartheta)) = \begin{cases} N_0(\vartheta) & \text{if } v = s \\ 0 & \text{if } v \in V \setminus \{s, t\}. \end{cases}$$

The functions x and ℓ are locally absolutely continuous. Due to the fundamental theorem of calculus for the Lebesgue integral, the families x and ℓ are defined by the initial conditions given by the empty network and their derivatives, which exist almost everywhere. The derivatives fulfill the following definition, which agrees with the definition of normalized thin flows with resetting by Cominetti, Correa, and Larré (2015) with the only difference that we do not fix the label of s to 1.

Definition 4.4 (Normalized thin flow with resetting). *For a shortest path graph with resetting $G' = (V, A', A^*)$, a static s - t flow $x' \in \mathbb{R}_{\geq 0}^{A'}$ in (V, A') is called a **normalized thin s - t flow with resetting** in G' if there exist labels $\ell' \in \mathbb{R}_{\geq 0}^V$ such that*

$$\begin{aligned} \ell'_w &= \min_{\substack{a \in \delta_{A'}^-(w) \\ a=(v,w)}} \varrho^a(\ell'_v, x'_a) \quad \text{for all } w \in V \setminus \{s\}, \text{ and} \\ \ell'_w &= \varrho^a(\ell'_v, x'_a) \quad \text{for all } a = (v, w) \in A' \text{ with } x'_a > 0, \end{aligned}$$

where the behavior of the labels along an arc $a = (v, w) \in A'$ is prescribed by the function

$$\varrho^a: \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}, \quad (\ell'_v, x'_a) \mapsto \begin{cases} \max\{\ell'_v, x'_a/\nu_a\} & \text{if } a \in A' \setminus A^* \\ x'_a/\nu_a & \text{if } a \in A^*. \end{cases}$$

Set $x' := dx/d\vartheta^+$ and $\ell' := d\ell/d\vartheta^+$ to be the right-derivatives of x and ℓ wherever they exist, which is almost everywhere. Differentiating the flow conservation constraints for $x(\vartheta)$ yields flow conservation for $x'(\vartheta)$. The Bellman equations (BE) and the complementarity conditions on a dynamic equilibrium imply that the flow $x'(\vartheta)$ is a normalized thin s - t flow of value $\nu_0(\vartheta)$ with resetting in G'_ϑ and has corresponding labels $\ell'(\vartheta)$. This characterization of the derivatives was found by Koch and Skutella (2011) for a slightly different notion of equilibria and investigated further by Cominetti, Correa, and Larré (2015).

The dynamic equilibrium from Figure 4.4 features four phases of different shortest path graphs with resetting G'_ϑ . They are displayed in Figure 4.5 together with corresponding thin flows with resetting.

4.3.1 A Linear Complementarity Problem

The complementarity conditions in Definition 4.4 make complementarity problems a natural candidate for describing normalized thin flows with resetting. Cominetti, Correa, and Larré (2011) take this approach with a nonlinear complementarity problem. The

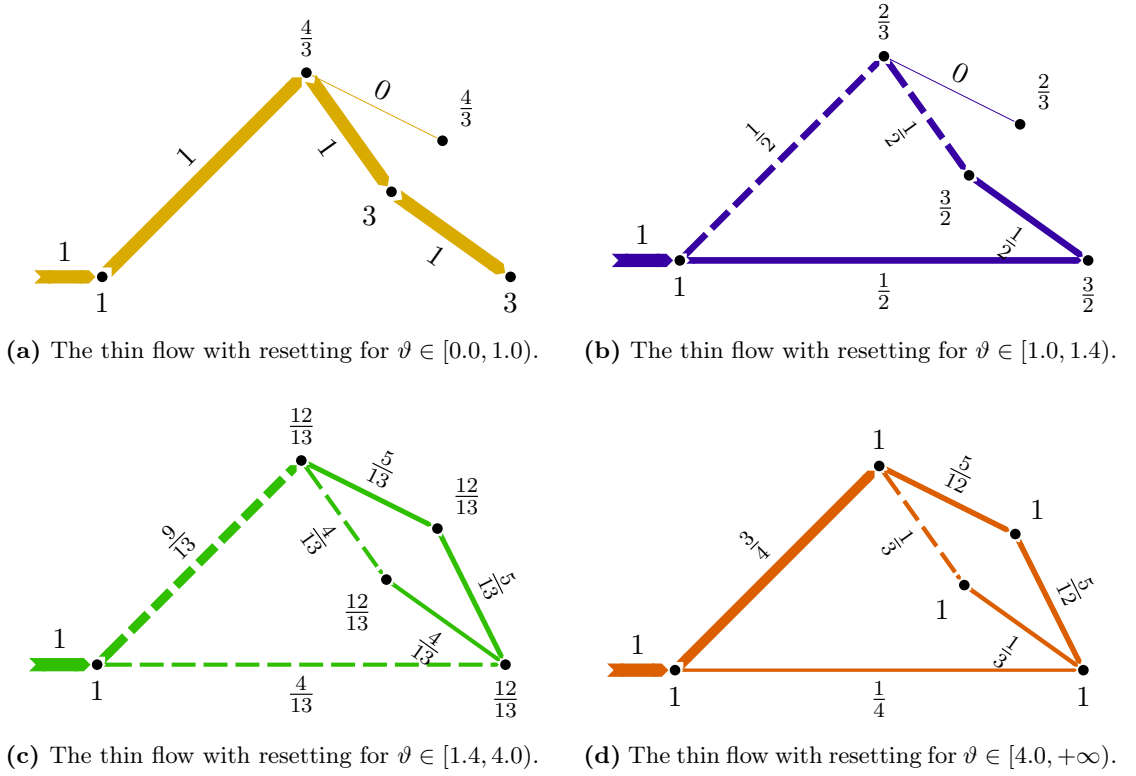


Figure 4.5: The thin flows with resetting in the four phases of the equilibrium depicted in Figure 4.4. The nonresetting and resetting arcs in the shortest path network are drawn solid and dashed, respectively. The vertices and arcs are annotated with the labels $\ell'(\vartheta)$ and the flow $x'(\vartheta)$, respectively.

nonlinearities in it are caused by taking minima and maxima. We observe that those can be resolved to linear complementarity conditions by introducing auxiliary variables.

Theorem 4.5 (Linear complementarity problem). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. Further, let $\nu'_0 \geq 0$ be an inflow rate and $\ell'_0 \geq 0$ a label. With the incidence matrix B of G' and the diagonal matrix $D := \text{diag}(\nu)$, define*

$$M := \begin{pmatrix} \mathbb{1}_s^\top & 0 & 0 \\ 0 & B_{V \setminus \{s\}} & 0 \\ -(B^+)^{\top} & D^{-1} & \text{Id}_{\bullet, A' \setminus A^*} \\ -(B_{\bullet, A' \setminus A^*}^-)^{\top} & D_{A' \setminus A^*, \bullet}^{-1} & \text{Id} \end{pmatrix}.$$

For $\ell' \in \mathbb{R}^V$, $x' \in \mathbb{R}^{A'}$, and $y' \in \mathbb{R}^{A' \setminus A^*}$, the following two statements are equivalent.

(i) $z' := (\ell', x', y')$ is a solution to the linear complementarity problem

$$z \geq 0, \quad Mz - \ell'_0 \mathbb{1}_s - \nu'_0 \mathbb{1}_t \geq 0, \quad z^\top (Mz - \ell'_0 \mathbb{1}_s - \nu'_0 \mathbb{1}_t) = 0 \quad (\text{LCP})$$

and additionally satisfies the **normalization constraints**

$$\ell'_w \geq \min_{v \in N^-(w)} \ell'_v \quad \text{for every } w \in V \setminus \{s\} \text{ with } \delta^-(w) \cap A^* = \emptyset.$$

(ii) ℓ' are the corresponding labels of the normalized thin s - t flow x' with resetting in G' of value ν'_0 and label $\ell'_s = \ell'_0$, and $y'_a = [\ell'_v - \frac{x'_a}{\nu'_a}]_+$ holds for all $a = (v, w) \in A' \setminus A^*$.

For the index sets of the rows and columns of the matrix M , we use the slight abuse of notation $V \cup A' \cup (A' \setminus A^*)$. The disjoint union of A' and $A' \setminus A^*$ refers to the fact that we regard $A' \setminus A^*$ as a copy of the respective subset of A' in this context. To indicate what row or column we refer to by an arc a , we explicitly distinguish between $a \in A'$ and $a \in A' \setminus A^*$. The former is associated with the variable x'_a , while the latter corresponds to y'_a .

Proof. Interpreting solutions as thin flows. Let $z' = (\ell', x', y')$ be a solution to (LCP) which satisfies $\ell'_w \geq \min_{v \in N^-(w)} \ell'_v$ for every $w \in V \setminus \{s\}$ with $\delta^-(w) \cap A^* = \emptyset$. The two complementary inequalities for ℓ'_s are $\ell'_s \geq 0$ and $\ell'_s \geq \ell'_0$. As $\ell'_0 \geq 0$, the latter has to be fulfilled with equality. For $a = (v, w) \in A' \setminus A^*$, the two inequalities associated with y'_a are $y'_a \geq 0$ and $y'_a \geq \ell'_v - \frac{x'_a}{\nu'_a}$. Together with the complementarity condition this yields $y'_a = [\ell'_v - \frac{x'_a}{\nu'_a}]_+$.

Let $w \in V$. For $a = (v, w) \in \delta^-(w)$, the inequality corresponding to x'_a yields

$$\begin{aligned} \ell'_w &\leq \frac{x'_a}{\nu'_a} &&= \varrho^a(\ell'_v, x'_a) && \text{if } a \in A^*, \text{ and} \\ \ell'_w &\leq \frac{x'_a}{\nu'_a} + y'_a = \max\{\ell'_v, \frac{x'_a}{\nu'_a}\} &&= \varrho^a(\ell'_v, x'_a) && \text{if } a \notin A^*. \end{aligned}$$

In both cases, the inequality is guaranteed to be tight if $x'_a > 0$ due to the complementarity conditions. It follows that $\ell'_w \leq \min_{a=(v,w) \in \delta^-(w)} \varrho^a(\ell'_v, x'_a)$ is valid and holds with equality in case that $\delta^-(w) \cap A^* \neq \emptyset$ or $\sum_{a \in \delta^-(w)} x'_a > 0$. Thanks to the normalization constraints, we also obtain equality in the remaining case which is $\delta^-(w) \cap A^* = \emptyset$ and $\sum_{a \in \delta^-(w)} x'_a = 0$.

We are left to show strict flow conservation for $w \in V \setminus \{s\}$. If $\ell'_w > 0$, it is implied by the complementarity condition for ℓ'_w . Thus, we assume $\ell'_w = 0$. For every $a = (v, w) \in \delta^-(w)$, we know from above that $x_a > 0$ would imply $0 = \ell'_w = \varrho^a(\ell'_v, x'_a)$ and, hence, the contradiction $x_a = 0$. Assuming $w = t$ and $\nu'_0 > 0$, the inequality corresponding to ℓ'_w yields the contradiction $0 \leq \sum_{a \in \delta^+(t)} x'_a \leq \sum_{a \in \delta^-(t)} x'_a - \nu'_0 < 0$. Otherwise, the inequality corresponding to ℓ'_w yields $\sum_{a \in \delta^+(w)} x'_a \leq \sum_{a \in \delta^-(w)} x'_a = 0$. It follows that $x'_a = 0$ for all $a \in \delta(w)$ which fulfills strict flow conservation at w .

In total, ℓ' proves x' to be a normalized thin s - t flow with resetting in G' of value ν'_0 with label $\ell'_s = \ell'_0$.

Interpreting thin flows as solutions. For the converse direction, let x' be a normalized thin s - t flow with resetting in G' of value ν'_0 and corresponding labels ℓ' with $\ell'_s = \ell'_0$. Set $y'_a := [\ell'_v - \frac{x'_a}{\nu'_a}]_+$ for all $a = (v, w) \in A' \setminus A^*$. We will show that $z' := (\ell', x', y')$ is

a solution to (LCP). z' clearly is nonnegative. The complementarity condition for ℓ'_s is fulfilled, as $M_{s,\bullet}z' = \ell'_s = \ell'_0$. Also, the complementarity conditions for $V \setminus \{s\}$ are met since x' is an s - t flow of value ν'_0 and, hence, $M_{V \setminus \{s\},\bullet}z' = B_{V \setminus \{s\}}x' = \nu'_0 \mathbb{1}_t$. Let $a = (v, w) \in A'$. The inequality associated with x'_a reads

$$\begin{aligned} \ell'_w &\leq \frac{x'_a}{\nu'_a} &&= \varrho^a(\ell'_v, x'_a) && \text{if } a \in A^*, \text{ and} \\ \ell'_w &\leq \frac{x'_a}{\nu'_a} + y'_a = \max\{\ell'_v, \frac{x'_a}{\nu'_a}\} &&= \varrho^a(\ell'_v, x'_a) && \text{if } a \notin A^*. \end{aligned}$$

In both cases, the respective inequality is valid due to the definition of normalized thin flows with resetting. Further, equality is guaranteed if $x'_a > 0$. This shows that the complementarity conditions are satisfied for the variable x'_a .

For $a = (v, w) \in A' \setminus A^*$, the inequality corresponding to variable y'_a reads $y'_a \geq \ell'_v - \frac{x'_a}{\nu'_a}$. This is clearly fulfilled by the above choice of y' . Further, $y'_a > 0$ implies $y'_a = [\ell'_v - \frac{x'_a}{\nu'_a}]_+ = \ell'_v - \frac{x'_a}{\nu'_a}$. Therefore, the complementarity condition for y'_a is satisfied as well.

The normalization constraints in (i) hold due to the definition of normalized thin flows with resetting. We conclude that (ii) implies (i). \blacksquare

Theorem 4.5 shows that a solution (ℓ', x', y') to (LCP) is nearly a normalized thin flows with resettings. Consulting its proof reveals that the normalization constraints in (i) are only needed for $w \in V \setminus \{s\}$ such that $\delta^-(w) \cap A^* = \emptyset$ and $\sum_{a \in \delta^-(w)} x'_a = 0$. In this case, (LCP) allows all values $0 \leq \ell'_w \leq \min_{v \in N^-(w)} \ell'_v$ while $\ell'_w = \min_{v \in N^-(w)} \ell'_v$ is necessary for ℓ' to be the corresponding labels of x' . As similarly observed by Cominetti, Correa, and Larré (2011) for their nonlinear complementarity problem, any solution to (LCP) can be normalized to fulfill (i). For that purpose, we define a normalization function $\pi: \mathbb{R}^V \rightarrow \mathbb{R}^V$. Let ℓ' be the labels of a solution to (LCP). Fix an arbitrary topological order of the vertices in G' and successively set for $w \in V$ in that order

$$\pi_w(\ell') = \begin{cases} \max\{\ell'_w, \min_{v \in N^-(w)} \pi_v(\ell')\} & \text{if } w \neq s \text{ and } \delta^-(w) \cap A^* = \emptyset \\ \ell'_w & \text{otherwise} \end{cases}.$$

Note that the definition is independent of the particular topological order that was chosen.

Lemma 4.6. *The normalization function $\pi: \mathbb{R}^V \rightarrow \mathbb{R}^V$ as defined above satisfies*

- (i) $\ell' \leq \pi(\ell')$ for all $\ell' \in \mathbb{R}_{\geq 0}^V$, and
- (ii) $\|\pi(\ell'') - \pi(\ell')\|_\infty \leq \|\ell'' - \ell'\|_\infty$ for all $\ell', \ell'' \in \mathbb{R}_{\geq 0}^V$.

Proof. (i) follows immediately from the definition of π . To see (ii), we prove by induction on $w \in V$ in topological order that $|\pi_w(\ell'') - \pi_w(\ell')| \leq \|\ell'' - \ell'\|_\infty$. For $w = s$ and all $w \in V$ such that $\delta^-(w) \cap A^* \neq \emptyset$, by definition $\pi_w(\ell') = \ell'_w$ and, hence, the statement holds. For $w \in V \setminus \{s\}$ with $\delta^-(w) \cap A^* = \emptyset$, assume that the induction hypothesis holds

for all vertices v which are topologically preceding w . Then, the claim also holds for w based on

$$|\pi_w(\ell'') - \pi_w(\ell')| \leq \max \left\{ |\ell''_w - \ell'_w|, \max_{v \in N^-(w)} |\pi_v(\ell'') - \pi_v(\ell')| \right\} \leq \|\ell'' - \ell'\|_\infty. \quad \blacksquare$$

Corollary 4.7 (Normalization). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{\geq 0}^{A'}$ be capacities. Further, let $\nu'_0 \geq 0$ be an inflow rate and $\ell'_0 \geq 0$ a label. If $z' = (\ell', x', y')$ is a solution to (LCP), then $\pi(\ell')$ are the corresponding labels of the normalized thin s - t flow x' with resetting in G' of value ν'_0 with label $\ell'_s = \ell'_0$.*

Proof. Let $z' = (\ell', x', y')$ be a solution to (LCP). Let $z'' = (\ell'', x'', y'')$ be another solution to (LCP) such that $x'' = x'$, $\ell'' \leq \pi(\ell')$, and the set $U := \{v \in V \mid \ell''_v < \pi_v(\ell')\}$ is inclusion-minimal. Note that z'' exists since Lemma 4.6 guarantees that z' is feasible to the optimization problem defining z'' . As seen in the proof of Theorem 4.5 for $a = (v, w) \in A$, the inequalities corresponding to x_a for z' and z'' read $\ell'_w \leq \varrho^a(\ell'_v, x'_a)$ and $\ell''_w \leq \varrho^a(\ell''_v, x''_a)$, respectively. Both hold with equality if $a \in A^*$ or $x'_a = x''_a > 0$. Assume $U \neq \emptyset$. From the definition of π , we obtain $s \notin U$. Let $w \in U$ be topologically minimal, in particular $N^-(w) \cap U = \emptyset$. Assume there is $a = (v, w) \in \delta^-(w)$ such that $a \in A^*$ or $x'_a > 0$. Then

$$\ell'_w = \varrho^a(\ell'_v, x'_a) \leq \varrho^a(\pi_v(\ell'), x'_a) = \varrho^a(\ell''_v, x''_a) = \ell''_w < \pi_w(\ell')$$

shows $a \notin A^*$. But then

$$\ell''_w < \pi_w(\ell') = \min_{u \in N^-(w)} \pi_u(\ell') = \min_{u \in N^-(w)} \ell''_u \leq \ell''_v = \varrho^a(\ell''_v, x''_a)$$

contradicts the complementarity requirements.

We conclude that $\delta^-(w) \cap A^* = \emptyset$ and x' vanishes on $\delta^-(w)$. Since weak flow conservation is part of (LCP), x' vanishes on $\delta^+(w)$ as well. Therefore, increasing ℓ''_w and y''_a for all $a \in \delta^+(w)$ to $\min_{v \in N^-(w)} \ell''_v$ yields a solution to (LCP) with strictly smaller U and contradicts the choice of z'' . Hence, $U = \emptyset$ and $\ell'' = \pi(\ell')$ have to hold. Consequently, $z'' = (\pi(\ell'), x', y'')$ also fulfills the normalization constraints in Theorem 4.5 (i) which implies the statement. \blacksquare

Corollary 4.7 allows us to use theory on linear complementarity problems to study normalized thin flows with resetting. In the literature, results on linear complementarity problems are typically stated in terms of the properties of their matrices. The next theorem regards the sign of the principal minors of the matrix M from (LCP). We need to clarify what is meant when speaking about the determinant of a square matrix without a linear ordering of its index set. Let I be a finite index set of size $k := |I|$ and $K \in \mathbb{R}^{I \times I}$ be a square matrix. For an arbitrary permutation $\sigma: [k] \rightarrow I$, define $\det(K) := \det(K_{\sigma(i), \sigma(j)})_{i, j \in [k]}$. Note that $\det(K)$ is independent of σ as the determinant is invariant under symmetric permutation of the rows and columns.

Lemma 4.8 (Principal minors of M). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. Further, let M be the corresponding matrix of (LCP). Then, all principal minors of M are nonnegative.*

The principal minors of M (in the algebraic sense) are closely related to the determinants of the Laplacian matrix of the minors of G' (in the graph theoretic sense). The Laplacian matrix appears naturally due to its relation to the incidence matrix as discussed in Section 2.2. Hence, the nonnegativity of the minors of M are ultimately based on the same property of the minors of Laplacian matrices.

Lemma 4.9 (Principal minors of a Laplacian). *All principal minors of a graph's weighted Laplacian matrix are nonnegative.*

Proof. Let $L \in \mathbb{R}^{V \times V}$ be the weighted Laplacian of a directed graph $G = (V, A')$ with weights $\nu \in \mathbb{R}_{>0}^{A'}$. For $U \subseteq V$, the principal submatrix $L_{U,U}$ is (weakly) column diagonally dominant, as for every $w \in U$

$$L_{w,w} = \nu(\delta_{A'}^-(w)) \geq \sum_{v \in N_{A'}^-(w) \cap U} \nu_{v,w} = \sum_{v \in U \setminus \{w\}} |L_{v,w}|.$$

It follows by the Geršgorin circle theorem Geršgorin 1931 that all real eigenvalues of $L_{U,U}$ are nonnegative and, hence, $\det(L_{U,U}) \geq 0$. \blacksquare

Proof of Lemma 4.8. Let the principal submatrix $M_{U \cup X \cup Y}$ of M be given by the index sets $U \subseteq V$, $X \subseteq A'$, and $Y \subseteq A' \setminus A^*$. We need to show $\det(M_{U \cup X \cup Y}) \geq 0$. If $s \in U$, a Laplacian expansion along the row of s shows $\det(M_{U \cup X \cup Y}) = \det(M_{U \setminus \{s\} \cup X \cup Y})$. Hence, we can assume that $s \notin U$.

Carving out the Laplacian. Refining the block structure of $M_{U \cup X \cup Y}$ by partitioning X into $X \setminus Y$ and $X \cap Y$ yields

$$\det(M_{U \cup X \cup Y}) = \det \begin{pmatrix} 0 & B_{U,X \setminus Y} & B_{U,X \cap Y} & 0 \\ -(B_{U,X \setminus Y}^+)^{\top} & D_{X \setminus Y, X \setminus Y}^{-1} & 0 & 0 \\ -(B_{U,X \cap Y}^+)^{\top} & 0 & D_{X \cap Y, X \cap Y}^{-1} & \text{Id}_{X \cap Y, Y} \\ -(B_{U,Y}^-)^{\top} & 0 & D_{Y, X \cap Y}^{-1} & \text{Id}_{Y, Y} \end{pmatrix}.$$

We multiply the rows in $X \setminus Y$ from the left by $B_{U,X \setminus Y} D_{X \setminus Y, X \setminus Y}^{-1}$ and subtract the result from the rows corresponding to U . Further, we subtract the rows in $X \cap Y \subseteq Y$ from the corresponding rows $X \cap Y \subseteq X$. As these are unitary row operations under which the determinant is invariant, it holds

$$\det(M_{U \cup X \cup Y}) = \det \begin{pmatrix} B_{U,X \setminus Y} D_{X \setminus Y, X \setminus Y}^{-1} (B_{U,X \setminus Y}^+)^{\top} & 0 & B_{U,X \cap Y} & 0 \\ -(B_{U,X \setminus Y}^+)^{\top} & D_{X \setminus Y, X \setminus Y}^{-1} & 0 & 0 \\ -(B_{U,X \cap Y}^+)^{\top} & 0 & 0 & 0 \\ -(B_{U,Y}^-)^{\top} & 0 & D_{Y, X \cap Y}^{-1} & \text{Id}_{Y, Y} \end{pmatrix}.$$

Symmetric rearranging of the rows and columns reveals the triangular block structure

$$\det(M_{U \cup X \cup Y}) = \det \begin{pmatrix} B_{U, X \setminus Y} D_{X \setminus Y, X \setminus Y} (B_{U, X \setminus Y}^+)^{\top} & B_{U, X \cap Y} & 0 & 0 \\ -(B_{U, X \cap Y})^{\top} & 0 & 0 & 0 \\ -(B_{U, Y}^-)^{\top} & D_{Y, X \cap Y}^{-1} & \text{Id}_{Y, Y} & 0 \\ -(B_{U, X \setminus Y}^+)^{\top} & 0 & 0 & D_{X \setminus Y, X \setminus Y}^{-1} \end{pmatrix}.$$

As discussed in [Section 2.2](#), the matrix $B_{\bullet, X \setminus Y} D_{X \setminus Y, X \setminus Y} (B_{\bullet, X \setminus Y}^+)^{\top}$ is the weighted Laplacian matrix of the graph $H := (V, X \setminus Y)$. Therefore, we will denote it by L^H . Applying the multiplicativity of the determinant for triangular block matrices results in

$$\det(M_{U \cup X \cup Y}) = \det(D_{X \setminus Y, X \setminus Y}^{-1}) \det \begin{pmatrix} L_{U, U}^H & B_{U, X \cap Y} \\ -(B_{U, X \cap Y})^{\top} & 0 \end{pmatrix}.$$

Dealing with incomplete arcs. If there exists $a \in (X \cap Y) \cap ((V \setminus U) \times (V \setminus U))$, that is, $a \in X \cap Y$ is not incident to any vertex in U , then $B_{U, a} = 0$. Therefore, $M_{U \cup X \cup Y}$ is singular in that case. On the other hand, assume there exists $a \in (X \cap Y) \cap \delta(U)$, that is, there is exactly one $u \in U$ which is incident to $a \in X \cap Y$. Applying the Laplacian expansion consecutively along the row a and column a yields

$$\det(M_{U \cup X \cup Y}) = \det(D_{X \setminus Y, X \setminus Y}^{-1}) B_{u, a}^2 \det \begin{pmatrix} L_{U \setminus \{u\}, U \setminus \{u\}}^H & B_{U \setminus \{u\}, X \cap Y} \\ -(B_{U \setminus \{u\}, X \cap Y})^{\top} & 0 \end{pmatrix}.$$

As $B_{u, a}^2 \geq 0$, we can assume that $X \cap Y \subseteq U \times U$ by using induction on $|U|$. Under this assumption, if $X \cap Y$ contains an undirected cycle, it follows that $B_{U, X \cap Y}$ does not have full column rank and $M_{U \cup X \cup Y}$ is singular. Thus, we additionally assume in the following that $X \cap Y$ does not contain any undirected cycle.

Contracting the nonresetting arcs. Let $U = \dot{\bigcup}_{i \in [k]} U_i$ and $X \cap Y = \dot{\bigcup}_{i \in [k]} Y_i$ give the partition of $(U, X \cap Y)$ into its $k \in \mathbb{N}$ weakly connected components $(U_1, Y_1), \dots, (U_k, Y_k)$. Then (U_i, Y_i) is a tree when ignoring the orientation of the arcs for every $i \in [k]$. Fix an arbitrary root vertex $r_i \in U_i$ for every $i \in [k]$ and define the set $R := \{r_1, \dots, r_k\}$. Refining the block structure of the remaining matrix by splitting U into $U \setminus R$ and R gives that

$$\det(M_{U \cup X \cup Y}) = \det(D_{X \setminus Y, X \setminus Y}^{-1}) \det \begin{pmatrix} L_{U \setminus R, U \setminus R}^H & L_{U \setminus R, R}^H & B_{U \setminus R, X \cap Y} \\ L_{R, U \setminus R}^H & L_{R, R}^H & B_{R, X \cap Y} \\ -B_{U \setminus R, X \cap Y}^{\top} & -B_{R, X \cap Y}^{\top} & 0 \end{pmatrix}.$$

For each $i \in [k]$ and $u \in U_i \setminus \{r_i\}$, we sequentially add row u to row r_i and column u to column r_i . Note that the determinant is invariant under these unitary operations and the resulting matrix does not depend on their order. The effect of the operations on $L_{U, U}^H$ can be interpreted as the contraction of arcs in the following sense. Let $\hat{H} := H/(X \cap Y)$ be

the graph that results from contracting all arcs of $X \cap Y$ in H , and $L^{\hat{H}}$ be its Laplacian matrix. For every $i \in [k]$, we identify r_i with the vertex which results from contracting U_i by Y_i . As (U_i, Y_i) is a weakly connected component with respect to $X \cap Y$, we obtain $\mathbb{1}_{U_i}^\top B_{U_i, X \cap Y} = 0$ for all $i \in [k]$. Putting this together gives

$$\det(M_{U \cup X \cup Y}) = \det\left(D_{X \setminus Y, X \setminus Y}^{-1}\right) \det \begin{pmatrix} L_{U \setminus R, U \setminus R}^H & * & B_{U \setminus R, X \cap Y} \\ * & L_{R, R}^{\hat{H}} & 0 \\ -B_{U \setminus R, X \cap Y}^\top & 0 & 0 \end{pmatrix},$$

where $*$ marks blocks the specific values of which are not relevant for our purposes. Since the components (U_i, Y_i) are trees, the matrix $B_{U \setminus R, X \cap Y}$ is square.

Concluding. By swapping the columns $U \setminus R$ with the columns $X \cap Y$, we obtain

$$\det(M_{U \cup X \cup Y}) = \det\left(D_{X \setminus Y, X \setminus Y}^{-1}\right) (-1)^{|X \cap Y|} \det \begin{pmatrix} B_{U \setminus R, X \cap Y} & * & L_{U \setminus R, U \setminus R}^H \\ 0 & L_{R, R}^{\hat{H}} & * \\ 0 & 0 & -B_{U \setminus R, X \cap Y}^\top \end{pmatrix}.$$

Exploiting the triangular block structure again and applying [Lemma 4.9](#) finally yields the statement

$$\det(M_{U \cup X \cup Y}) = \det\left(D_{X \setminus Y, X \setminus Y}^{-1}\right) \det(B_{U \setminus R, X \cap Y})^2 \det\left(L_{R, R}^{\hat{H}}\right) \geq 0. \quad \blacksquare$$

From [Lemma 4.8](#), we obtain a proof for the existence of normalized thin s - t flows with resetting. It is an alternative to the original proof by Cominetti, Correa, and Larré (2015), which is based on an elegant application of Kakutani's fixed point theorem.

Theorem 4.10 (Existence). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and let $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. For all given values $\nu'_0, \ell'_0 \geq 0$, there exists a normalized thin s - t flow with resetting of value ν'_0 in G' and corresponding label ℓ'_0 at s .*

Proof. Due to a result by Cottle et al. (2009, Theorem 3.9.22) and [Corollary 4.7](#), it is enough to show that $z' = 0$ is the unique solution to (LCP) for $\nu'_0 = 0$ and $\ell'_0 = 0$, and that its matrix M has only nonnegative minors. The latter is taken care of by [Lemma 4.8](#). For the flow rate $\nu'_0 = 0$ and label $\ell'_0 = 0$, the unique s - t flow $x' \equiv 0$ is also a normalized thin s - t flow with resetting in G' with unique corresponding labels $\ell' \equiv 0$. By [Theorem 4.5](#) and [Corollary 4.7](#), $z' = 0$ is therefore the unique solution to (LCP) for these parameters. \blacksquare

The proof of [Theorem 4.10](#) uses the uniqueness of the solution $z' = 0$ to (LCP) for $\nu'_0 = 0$ and $\ell'_0 = 0$. Together with [Lemma 4.8](#), it even proves the applicability of known pivoting methods and iterative methods for linear complementarity problems to (LCP). In particular, a result from Cottle et al. (2009, Theorems 4.4.8 and 4.4.11) shows that Lemke's algorithm can be used to find a normalized thin flow with resetting in finitely many steps, when dealing with degeneracy appropriately (see Cottle et al. 2009, Section 4.9).

4.3.2 Parametric Normalized Thin Flows with Resetting

We want to examine the dependency of normalized thin flows with resetting on the flow value and the label of s . The following proof for the monotonicity of the corresponding labels in these two parameters is a refinement of the analysis that Cominetti, Correa, and Larré (2015) use to show uniqueness of the labels (which is an immediate corollary of the monotonicity).

Theorem 4.11 (Monotonicity of labels). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. Further, let x', x'' be two normalized thin s - t flows with resetting in G' and let ℓ', ℓ'' be corresponding labels, respectively. If the flow values fulfill $|x'| \leq |x''|$ and the labels fulfill $\ell'_s \leq \ell''_s$, then $\ell' \leq \ell''$ holds (element-wise).*

Proof. Assume that $\ell'_s \leq \ell''_s$ and $|x'| \leq |x''|$ hold. We examine the thin flow with resetting across the cut $U := \{v \in V \mid \ell'_v > \ell''_v\}$ and conclude that it has to be empty. Assume for a contradiction that $U \neq \emptyset$.

We claim that $x'_a = x''_a$ for all $a \in \delta_{A'}(U)$. Assume again this is wrong. We know $s \notin U$ and, thus,

$$\sum_{a \in \delta_{A'}^+(U)} (x'_a - x''_a) - \sum_{a \in \delta_{A'}^-(U)} (x'_a - x''_a) = \sum_{u \in U} \left(\left(\sum_{a \in \delta_{A'}^+(u)} x'_a - \sum_{a \in \delta_{A'}^-(u)} x'_a \right) - \left(\sum_{a \in \delta_{A'}^+(u)} x''_a - \sum_{a \in \delta_{A'}^-(u)} x''_a \right) \right) \geq 0.$$

Therefore, there has to be $a = (v, w) \in A'$ such that $a \in \delta_{A'}^+(U)$ and $x'_a > x''_a$, or $a \in \delta_{A'}^-(U)$ and $x'_a < x''_a$. If $a \in \delta_{A'}^+(U)$ then $\ell'_w = \varrho^a(\ell'_v, x'_a) > \varrho^a(\ell''_v, x''_a) \geq \ell''_w$ which contradicts $w \notin U$. If $a \in \delta_{A'}^-(U)$ then $\ell''_w = \varrho^a(\ell''_v, x''_a) \geq \varrho^a(\ell'_v, x'_a) \geq \ell'_w$ which contradicts $w \in U$. Thus, x' and x'' agree on $\delta_{A'}(U)$.

As a consequence, $\delta_{A'}^-(U) \cap A^* = \emptyset$ and $\varrho^a(\ell'_v, x'_a) = \ell'_v$ for all $a = (v, w) \in \delta_{A'}^-(U)$. Since A' is acyclic and $s \notin U$, there is $w \in U$ such that $\emptyset \neq \delta_{A'}^-(w) \subseteq \delta_{A'}^-(U)$. Then, $w \in U$ contradicts

$$\ell'_w = \min_{\substack{a \in \delta_{A'}^-(w) \\ a=(v,w)}} \varrho^a(\ell'_v, x'_a) = \min_{v \in N^-(w)} \ell'_v \leq \min_{v \in N^-(w)} \ell''_v \leq \min_{\substack{a \in \delta_{A'}^-(w) \\ a=(v,w)}} \varrho^a(\ell''_v, x''_a) = \ell''_w. \quad \blacksquare$$

Corollary 4.12 (Uniqueness of labels). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. The corresponding labels ℓ' of a normalized thin s - t flows x' with resetting in G' are uniquely determined by its flow value $|x'|$ and the label ℓ'_s .*

Example 4.13 (Nonuniqueness of thin flows). Consider the shortest path graph with resetting $G' = (V, A', A^*)$ consisting of the source s , the sink t , and two parallel nonresetting arcs a and b from s to t , that is, $V = \{s, t\}$, $A' = \{a, b\}$ and $A^* = \emptyset$. Note that the network can be transformed to an equivalent instance without multi-arcs by introducing additional vertices on arcs. Let $\nu \equiv 1$. Then for every $0 \leq x'_a \leq 1$ setting $x'_b = 1 - x'_a$ defines a normalized thin s - t flow with resetting of value 1 in G' with corresponding labels $\ell'_s = \ell'_t = 1$.

Corollary 4.12 shows that the corresponding labels ℓ' of a normalized thin s - t flow with resetting are uniquely determined by the label ℓ'_s and the flow rate ν'_0 . In contrast to that, **Example 4.13** shows that the flow itself is not necessarily uniquely determined by those parameters. In some sense, however, its nonuniqueness is the only kind that can occur. The flow on a subset of arcs is uniquely determined by the labels. For every $a = (v, w) \in A' \setminus A^*$ with $\ell'_v < \ell'_w$ and for every $a = (v, w) \in A^*$, it holds $x'_a = \ell'_w \nu_a$. For every $a = (v, w) \in A' \setminus A^*$ with $\ell'_v > \ell'_w$ on the other hand, $x'_a = 0$. Therefore, nonuniqueness can only arise in weakly connected components with respect to $A' \setminus A^*$ of constant label.

The proof of **Corollary 4.12** is combinatorial in nature and it is not clear whether a similar result can be obtained by the means of linear complementarity problems. There are results on the uniqueness of the solutions to linear complementarity problems. Those, however, do not apply to **(LCP)**. Apart from the discussed nonuniqueness of **Example 4.13**, **(LCP)** suffers from a different kind which is the nonnormalized labels. The latter can be addressed as discussed at the end of this section. Yet, it would only yield uniqueness of the labels. To the knowledge of the author, there are no results on such partial uniqueness in the general theory on linear complementarity problems.

The flow value ν'_0 appears in a linear way in **(LCP)**. This allows to treat it as a variable and get a new linear complementarity problem. Its set of solutions captures the dependency of normalized thin flows with resetting on the flow value, which is analyzed in the following lemmas and used in **Sections 4.4** and **4.5**.

Proposition 4.14 (Parametric thin flows with resetting). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting and $\nu \in \mathbb{R}_{>0}^{A'}$ be capacities. Then there are continuous, piecewise linear functions $\varrho^{G'} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}^V$ and $\chi^{G'} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}^{A'}$ such that, for all $\nu'_0 \geq 0$, the vector $\chi^{G'}(\nu'_0)$ is a normalized thin s - t flow with resetting of value ν'_0 in G' and $\varrho^{G'}(\nu'_0)$ are corresponding labels with $\varrho_s^{G'}(\nu'_0) = 1$.*

Proof. The labels. Define the function $\varrho^{G'} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}^V$ by setting, for every $\nu'_0 \geq 0$, the vector $\varrho^{G'}(\nu'_0)$ to be the corresponding labels of a normalized thin s - t flow with resetting of value ν'_0 in G' such that the corresponding label at s is 1. By **Theorem 4.10** and **Corollary 4.12**, this is a sound definition. According to Cottle, Pang, and Stone (2009, Theorem 7.2.1), the set of solutions to a linear complementarity problem depends in a locally upper Lipschitz continuous way on its right-hand side as follows. Fix $\nu'_0 \geq 0$. Let $z'' = (\ell'', x'', y'')$ be a solution to **(LCP)** with ν'_0 replaced by another flow rate $\nu''_0 \geq 0$. There exists a constant $C > 0$ and $\varepsilon > 0$ such that $|\nu''_0 - \nu'_0| \leq \varepsilon$ implies the existence of a solution $z' = (\ell', x', y')$ to **(LCP)** satisfying $\|z'' - z'\|_\infty \leq C|\nu''_0 - \nu'_0|$. Together with **Lemma 4.6** and **Corollary 4.7**, it follows

$$\|\varrho^{G'}(\nu''_0) - \varrho^{G'}(\nu'_0)\|_\infty = \|\pi(\ell'') - \pi(\ell')\|_\infty \leq \|\ell'' - \ell'\|_\infty \leq \|z'' - z'\|_\infty \leq C|\nu''_0 - \nu'_0|,$$

that is, $\varrho^{G'}$ is locally Lipschitz continuous. In order to see piecewise linearity, we consider a slight variant of (LCP) where ν'_0 is regarded a variable and ℓ'_0 is fixed to one, that is,

$$z \geq 0, \quad \begin{pmatrix} 0 & 0 \\ -\mathbb{1}_t & M \end{pmatrix} z - \mathbb{1}_s \geq 0, \quad z^\top \left(\begin{pmatrix} 0 & 0 \\ -\mathbb{1}_t & M \end{pmatrix} z - \mathbb{1}_s \right) = 0. \quad (\text{LCP}')$$

Its set of solutions corresponds exactly to all sets of solutions to (LCP) with arbitrary flow rate $\nu'_0 \geq 0$ and $\ell'_0 = 1$. Then using again Lemma 4.6 and Corollary 4.7, the hypograph of $\varrho^{G'}$ can be written as

$$\begin{aligned} \text{hyp}(\varrho^{G'}) &= \text{graph}(\varrho^{G'}) + (\{0\} \times \mathbb{R}_{\leq 0}^V) \\ &= \left\{ (\nu'_0, \varrho^{G'}(\nu'_0)) \in \mathbb{R} \times \mathbb{R}^V \mid \nu'_0 \geq 0 \right\} + (\{0\} \times \mathbb{R}_{\leq 0}^V) \\ &= \left\{ (\nu'_0, \pi(\ell')) \in \mathbb{R} \times \mathbb{R}^V \mid (\nu'_0, \ell', x', y') \text{ solves } (\text{LCP}') \right\} + (\{0\} \times \mathbb{R}_{\leq 0}^V) \\ &= \left\{ (\nu'_0, \ell') \in \mathbb{R} \times \mathbb{R}^V \mid (\nu'_0, \ell', x', y') \text{ solves } (\text{LCP}') \right\} + (\{0\} \times \mathbb{R}_{\leq 0}^V) \end{aligned}$$

By Cottle, Pang, and Stone (2009, p. 646), the first (Minkowski) summand of the right-hand side is the union of finitely many polyhedra as it is the linear projection of the set of solutions to a linear complementarity problem. Hence, the same holds for $\text{hyp}(\varrho^{G'})$ and $\text{graph}(\varrho^{G'})$. This yields that $\varrho^{G'}$ must be piecewise linear.

The flow. We would like to define $\chi^{G'} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}^{A'}$ in a similar way as $\varrho^{G'}$ such that, for every $\nu'_0 \geq 0$, the vector $\chi^{G'}(\nu'_0)$ is a normalized thin s - t flow with resetting of value ν'_0 in G' that admits the corresponding label 1 at s . As Example 4.13 shows, the flow values on the arcs are generally not unique. $\chi^{G'}$ can be chosen in any way such that its graph lies within the projection of the set of solutions to (LCP') onto the flow rate ν'_0 and flow variables x' . Reasoning as above, it follows that this can be done such that $\chi^{G'}$ is piecewise linear and continuous as well. ■

Lemma 4.15 (Homogeneity and monotonicity). *Let $G' = (V, A', A^*)$ be a shortest path graph with resetting.*

- (i) *For all $\nu'_0, \ell'_0 \geq 0$, the vector $\ell'_0 \cdot \varrho^{G'}(\nu'_0/\ell'_0)$ are the corresponding labels of the normalized thin s - t flow $\ell'_0 \cdot \chi^{G'}(\nu'_0/\ell'_0)$ with resetting in G' of value ν'_0 .*
- (ii) *For all $v \in V$, the function $\varrho_v^{G'}$ is nondecreasing and $\nu'_0 \mapsto 1/\nu'_0 \cdot \varrho_v^{G'}(\nu'_0)$ is nonincreasing.*

Proof. Consulting Theorem 4.5, it becomes quite obvious that a normalized thin flow with resetting with corresponding label 1 at s and flow value ν'_0/ℓ'_0 can be scaled by ℓ'_0 to yield a normalized thin flow with resetting with corresponding label ℓ'_0 at s and flow value ν'_0 . This yields the first statement.

The monotonicity of $\varrho^{G'}$ is implied directly by [Theorem 4.11](#). By the above, the scaled vector $1/\nu'_0 \cdot \varrho^{G'}(\nu'_0)$ can be interpreted as the corresponding labels of a normalized thin flow of value 1 with label $1/\nu'_0$ at s . Its monotonicity follows again from [Theorem 4.11](#). ■

We finish this section by arguing that the normalization constraints of [Theorem 4.5 \(i\)](#) can be incorporated in (LCP). The resulting linear complementarity problem is an exact formulation for normalized thin flows with resetting. This extension, however, requires a rather technical construction, which is why the presentation with successive normalization was chosen. In order to model the normalization constraints for $w \in V$ with $\delta^-(w) \cap A^* = \emptyset$, fix an arbitrary ordering $N^-(w) = \{v_1, \dots, v_n\}$ where $n = |N^-(w)|$. Introduce new variables $d_{w,i}$ for $i \in [n]$ and add the complementarity conditions

$$\begin{aligned} d_{w,1} &\geq 0, & \ell'_w - \ell'_{v_n} + d_{w,n} &\geq 0, & d_{w,1}(\ell'_w - \ell'_{v_n} + d_{w,n}) &= 0 \\ d_{w,2} &\geq 0, & \ell'_{v_1} - \ell'_{v_2} + d_{w,2} &\geq 0, & d_{w,2}(\ell'_{v_1} - \ell'_{v_2} + d_{w,2}) &= 0 \\ d_{w,i} &\geq 0, & \ell'_{v_{i-1}} - d_{w,i-1} - \ell'_{v_i} + d_{w,i} &\geq 0, & d_{w,i}(\ell'_{v_{i-1}} - d_{w,i-1} - \ell'_{v_i} + d_{w,i}) &= 0 \end{aligned}$$

for all $3 \leq i \leq n$. It can be shown by induction that these conditions are equivalent to $\ell'_{v_i} - d_{w,i} = \min_{j \in [i]} \ell'_{v_j}$ for all $i \in [n]$ and $\ell'_w \geq \ell'_{v_n} - d_{w,n} = \min_{j \in [n]} \ell'_{v_j} = \min_{v \in N^-(w)} \ell'_v$.

4.4 Evolution of Dynamic Equilibria

We extend the constructive method for dynamic equilibria by Koch and Skutella ([2011](#)) to inflow rates $\nu_0 \in L^1_{loc}(\mathbb{R})$ which are right-monotone (in addition to being nonnegative and vanishing almost everywhere on $\mathbb{R}_{<0}$). A definition of right-monotonicity follows the next theorem, which holds for arbitrary inflow rates.

4.4.1 A Differential Equation

Consider the earliest times function $\ell: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^V$ of a dynamic equilibrium. For $\vartheta \in \mathbb{R}$, the sets of active and resetting arcs agree with A^{ℓ}_{ϑ} and A^*_{ϑ} and, therefore, are determined by $\ell(\vartheta)$ only. If ℓ is right-differentiable at $\vartheta \in \mathbb{R}$, its right-derivative $d\ell/d\vartheta^+$ at ϑ represents the corresponding labels of a normalized thin flow with resetting of value $\nu_0(\vartheta)$ in G'_{ϑ} . Hence, also $d\ell/d\vartheta^+(\vartheta)$ is determined by $\ell(\vartheta)$. Indeed, the earliest times function of dynamic equilibria can be characterized as the set of solutions to a differential equation, as the next theorem states.

In the case that ℓ is the earliest times function of a dynamic equilibrium, we know that G'_{ϑ} is a shortest path graph with resetting for all $\vartheta \in \mathbb{R}$. To ensure this property through the differential equation, we need to extend the definition of $\varrho^{G'}$ and $\chi^{G'}$ to triples $G' = (V, A', A^*)$ such that $A^* \subseteq A' \subseteq A$ where A' is acyclic, but there are $v \in V$ without an s - v path in A' . In that case, let $U := \{v \in V \mid \exists s$ - v path in $A'\}$, set $\varrho^{G'}_{v'} \equiv 0$ for $v \in U$, and $\varrho^{G'}_{v'} \equiv 1$ for $v \in V \setminus U$. Further, define $\chi^{G'}_{a'} \equiv 0$ for all $a \in A'$.

Theorem 4.16 (Differential equation). *Let ℓ_v^0 denote the shortest distance from s to v in G with respect to τ for all $v \in V$. Then $\ell: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^V$ is the earliest times function of a dynamic equilibrium if and only if ℓ is a locally absolutely continuous solution to*

$$\ell(\vartheta) = (\vartheta + \ell_v^0)_{v \in V} \text{ for all } \vartheta \leq 0 \quad \text{and} \quad \frac{d\ell}{d\vartheta}(\vartheta) = \varrho^{G'_\vartheta}(\nu_0(\vartheta)) \text{ for a.e. } \vartheta \geq 0. \quad (\text{DE})$$

Proof. Assume ℓ is the earliest times function of a dynamic equilibrium. Since the network is assumed to be empty up to time zero, the initial condition on $\mathbb{R}_{\leq 0}$ holds due to (BE). The components of ℓ are locally absolutely continuous and, hence, differentiable almost everywhere. Cominetti, Correa, and Larré (2015) show that, if it exists, the right-derivative of ℓ at time $\vartheta \in \mathbb{R}$ represents the corresponding labels of a normalized thin flow with resetting of value $\nu_0(\vartheta)$ in G'_ϑ with label 1 at s . Hence, $\frac{d\ell}{d\vartheta}(\vartheta) = \varrho^{G'_\vartheta}(\nu_0(\vartheta))$ follows for almost every $\vartheta \geq 0$.

For the converse direction, assume ℓ is a locally absolutely continuous solution to the differential equation (DE). ℓ defines $G'_\vartheta = (V, A'_\vartheta, A_\vartheta^*)$ for all $\vartheta \in \mathbb{R}$. To simplify notation, we set

$$\varrho: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^V, \vartheta \mapsto \varrho^{G'_\vartheta}(\nu_0(\vartheta)) \quad \text{and} \quad \chi: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}^A, \vartheta \mapsto \begin{cases} \chi_a^{G'_\vartheta}(\nu_0(\vartheta)) & \text{for } a \in A'_\vartheta \\ 0 & \text{for } a \notin A'_\vartheta. \end{cases}$$

Note that in both definitions the shortest path graph G'_ϑ as well as the inflow rate $\nu_0(\vartheta)$ depend on the parameter ϑ . Moreover, the codomain of $\chi^{G'_\vartheta}$ depends on ϑ . χ extends the definition of the flows on A'_ϑ onto A by zero. The fundamental theorem of calculus yields

$$\ell(\vartheta) = \ell^0 + \int_0^\vartheta \varrho(\theta) d\theta \quad \text{for all } \vartheta \in \mathbb{R},$$

where integration is applied element-wise. Since $\varrho_v \geq 0$ for all $v \in V$, ℓ_v is a monotonically nondecreasing function. Based on ℓ , we define a flow over time and prove step by step that it is feasible, has the earliest times function ℓ , and satisfies the conditions of a dynamic equilibrium.

The shortest path graphs with resetting. We start by showing that G'_ϑ is a shortest path graph with resetting for all $\vartheta \in \mathbb{R}$. It is clear from the initial condition that the claim holds for $\vartheta \leq 0$. Assume for contradiction, that there is $\vartheta > 0$ such that A'_ϑ is not a shortest path graph with resetting. This means that the set $U := \{v \in V \mid \exists s\text{-}v \text{ path in } A'_\vartheta\}$ is a proper subset of V and $\delta^+(U) \cap A'_\vartheta = \emptyset$. Let $\vartheta := \sup\{\theta \leq \vartheta \mid \delta^+(U) \cap A'_\theta \neq \emptyset\}$ be the last time before ϑ that an arc leaving U was active. Then, $0 \leq \vartheta < \vartheta$ due to the initial condition and the continuity of ℓ . In particular, it holds $\delta^+(U) \cap A'_\vartheta \neq \emptyset$, and $\delta^+(U) \cap A'_\theta = \emptyset$ for all $\theta \in (\vartheta, \vartheta]$. Then for all $(v, w) \in \delta^+(U) \cap A'_\vartheta$, $\varrho_w \equiv 1$ and $\varrho_w - \varrho_v \geq 0$ on $(\vartheta, \vartheta]$ yield $\emptyset \neq \delta^+(U) \cap A'_\vartheta \subseteq \delta^+(U) \cap A'_\vartheta = \emptyset$, a contradiction.

As an immediate consequence, $\lim_{\vartheta \rightarrow +\infty} \ell_v(\vartheta) \geq \lim_{\vartheta \rightarrow +\infty} \ell_s(\vartheta) + \ell_v^0 = +\infty$ holds for all $v \in V$. On the other hand, $\lim_{\vartheta \rightarrow -\infty} \ell_v(\vartheta) = \lim_{\vartheta \rightarrow -\infty} \vartheta + \ell_v^0 = -\infty$. The continuity of ℓ yields $\ell_v(\mathbb{R}) = \mathbb{R}$ for all $v \in V$.

The flow over time. For $a = (v, w) \in A$, define the flow functions $f_a^+ : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ and $f_a^- : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ by setting

$$(f_a^+ \circ \ell_v) \cdot \varrho_v \equiv \chi_a \equiv (f_a^- \circ \ell_w) \cdot \varrho_w \quad \text{a.e. on } \mathbb{R}.$$

We shall see that this is a sound definition. For $\vartheta < \bar{\vartheta}$ such that $\ell_v(\vartheta) = \ell_v(\bar{\vartheta})$, the monotonicity of ℓ_v implies $\ell_v \equiv \ell_v(\vartheta)$ and $\varrho_v \equiv 0$ almost everywhere on $[\vartheta, \bar{\vartheta}]$. Hence, f_a^+ is well-defined at almost every $\ell_v(\vartheta)$ for $\vartheta \in \mathbb{R}$ such that $d\ell_v/d\vartheta(\vartheta) = \varrho_v(\vartheta) \neq 0$. As the set $\{\ell_v(\vartheta) \mid \vartheta \in \mathbb{R} \text{ and } d\ell_v/d\vartheta(\vartheta) = 0\}$ has measure zero, f_a^+ is well-defined almost everywhere on $\ell_v(\mathbb{R}) = \mathbb{R}$. Reasoning in a similar way gives that f_a^- is also well-defined almost everywhere on \mathbb{R} .

Definition 4.4 yields almost everywhere $f_a^- \leq \nu_a$. Similarly, almost everywhere $f_a^+ \equiv f_a^+ \circ \ell_s \equiv \chi_a / \varrho_s \leq \nu_0$ if $v = s$ and $f_a^+ \leq \nu(\delta^-(v))$ if $v \neq s$. As ν_0 is locally integrable and $(\nu_b)_{b \in A}$ is constant, the functions f_a^+ and f_a^- are locally integrable as well. Further, f_a^+ and f_a^- vanish on $\mathbb{R}_{\leq 0}$ as ν_0 does so and $\ell(0) = \ell^0 \geq 0$. The flow conservation constraints hold, since normalized thin flows with resetting obey them.

The queue lengths. For $a = (v, w) \in A$, we define the function z_a via its derivative and show that it evaluates to the queue length of a as induced by f_a^+ . For that purpose, we define $dz_a/d\vartheta$ by setting

$$\frac{dz_a}{d\vartheta}(\ell_v(\vartheta)) \cdot \varrho_v(\vartheta) = \begin{cases} 0 & \text{if } a \in A \setminus A'_\vartheta \\ \nu_a \cdot [\varrho_w(\vartheta) - \varrho_v(\vartheta)]_+ & \text{if } a \in A'_\vartheta \setminus A^*_\vartheta \\ \nu_a \cdot (\varrho_w(\vartheta) - \varrho_v(\vartheta)) & \text{if } a \in A^*_\vartheta \end{cases} \quad \text{for a.e. } \vartheta \in \mathbb{R}.$$

As argued similarly for the flow functions, this determines $dz_a/d\vartheta$ almost everywhere on \mathbb{R} as the set $\{\ell_v(\vartheta) \mid \vartheta \in \mathbb{R} \text{ and } d\ell_v/d\vartheta(\vartheta) = 0\}$ has measure zero. By using that $0 \leq \varrho_w \leq \max\{\varrho_v, \chi_a/\nu_a\}$, we can bound $dz_a/d\vartheta$ almost everywhere on $\ell_v(\mathbb{R}) = \mathbb{R}$ by noting that for almost every $\vartheta \in \mathbb{R}$

$$\begin{aligned} \left| \frac{dz_a}{d\vartheta}(\ell_v(\vartheta)) \right| \cdot \varrho_v(\vartheta) &\leq \nu_a \cdot |\varrho_w(\vartheta) - \varrho_v(\vartheta)| \\ &\leq \max\{\nu_a \cdot \varrho_v(\vartheta), \chi_a(\vartheta)\} = \max\{\nu_a, f_a^+(\ell_v(\vartheta))\} \cdot \varrho_v(\vartheta). \end{aligned}$$

This shows that the local integrability of f_a^+ carries over to $dz_a/d\vartheta$. Therefore, $dz_a/d\vartheta$ is indeed the derivative of the function $z_a : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}, \vartheta \mapsto \int_{\ell_v(0)}^{\vartheta} dz_a/d\vartheta(\theta) d\theta$. Since ℓ_v, ℓ_w , and z_a are locally absolutely continuous and ℓ_v is monotone, the formula for the change of variables applies to $z_a \circ \ell_v$ and yields

$$\begin{aligned} z_a(\ell_v(\vartheta)) &= \int_0^{\vartheta} \frac{dz_a}{d\vartheta}(\ell_v(\theta)) \cdot \varrho_v(\theta) d\theta \\ &= \int_0^{\vartheta} \nu_a \cdot \frac{d}{d\theta} [\ell_w(\theta) - \ell_v(\theta) - \tau_a]_+ d\theta = \nu_a \cdot [\ell_w(\vartheta) - \ell_v(\vartheta) - \tau_a]_+. \end{aligned}$$

Now, we can show that z_a is indeed the queue length that is induced by f_a^+ by proving that z_a is a solution to (QD). The above identity implies $z_a \equiv 0$ on $\mathbb{R}_{\leq 0} \subseteq (-\infty, \ell_v^0]$. Let $\vartheta \in \mathbb{R}$ such that $\varrho_v(\vartheta) \neq 0$. In the case $a \in A \setminus A'_\vartheta$, we know from the above identity and continuity of ℓ that $z_a(\ell_v(\vartheta)) = 0$ and $dz_a/d\vartheta(\ell_v(\vartheta)) = 0$. On the other hand, the definition of χ yields $\chi_a(\vartheta) = 0$, which implies $[f_a^+(\ell_v(\vartheta)) - \nu_a]_+ = 0$ and, thus, (QD) holds for a at time $\ell_v(\vartheta)$. If $a \in A'_\vartheta \setminus A_\vartheta^*$, we also know $z_a(\ell_v(\vartheta)) = 0$. In the case that $\nu_a \cdot \varrho_w(\vartheta) = \max\{\nu_a \cdot \varrho_v(\vartheta), \chi_a(\vartheta)\}$, it follows immediately that $dz_a/d\vartheta(\ell_v(\vartheta)) = [f_a^+(\ell_v(\vartheta)) - \nu_a]_+$. Otherwise, $\chi_a(\vartheta) = 0$ must hold, which implies the equation $dz_a/d\vartheta(\ell_v(\vartheta)) = 0 = [f_a^+(\ell_v(\vartheta)) - \nu_a]_+$. Finally, if $a \in A_\vartheta^*$, then $z_a(\ell_v(\vartheta)) > 0$ holds. Further, $\nu_a \cdot \varrho_w(\vartheta) = \chi_a(\vartheta)$ yields $dz_a/d\vartheta(\ell_v(\vartheta)) = f_a^+(\ell_v(\vartheta)) - \nu_a$. In total, z_a fulfills (QD) for almost every $\ell_v(\vartheta)$ such that $\vartheta \in \mathbb{R}$ and $\varrho_v(\vartheta) \neq 0$, which is almost everywhere on \mathbb{R} .

The queuing dynamics. To see that (f^+, f^-) is a feasible flow over time, we will show that it respects the queuing dynamics $z_a(\ell_v(\vartheta)) = F_a^+(\ell_v(\vartheta)) - F_a^-(\ell_v(\vartheta) + \tau_a)$ for all $a \in A$ and $\vartheta \in \mathbb{R}$.

Let $a = (v, w) \in A$. We start by establishing it for a single point in time and extend it by looking at its derivative. Consider the time $\hat{\vartheta} = \inf\{\theta \geq 0 \mid a \in A'_\theta\}$, relative to which a is active for the first time. If no such time exists, that is, $\hat{\vartheta} = +\infty$, then $z_a \equiv 0$, $F_a^+ \equiv 0$, and $F_a^- \equiv 0$ satisfy the equation trivially. Otherwise, we get $\ell_w(\hat{\vartheta}) = \ell_v(\hat{\vartheta}) + \tau_a$ and, therefore,

$$z_a(\ell_v(\hat{\vartheta})) = 0 = F_a^+(\ell_v(\hat{\vartheta})) - F_a^-(\ell_w(\hat{\vartheta})) = F_a^+(\ell_v(\hat{\vartheta})) - F_a^-(\ell_v(\hat{\vartheta}) + \tau_a).$$

It remains to prove that the derivatives of both sides of the equation agree almost everywhere. As we have shown that z_a is a solution to (QD), it suffices to prove for almost every $\vartheta \in \mathbb{R}$ with $\varrho_v(\vartheta) \neq 0$ that

$$f_a^-(\ell_v(\vartheta) + \tau_a) = \begin{cases} \min\{f_a^+(\ell_v(\vartheta)), \nu_a\} & \text{if } z_a(\ell_v(\vartheta)) = 0, \text{ that is, } a \in A \setminus A_\vartheta^* \\ \nu_a & \text{if } z_a(\ell_v(\vartheta)) > 0, \text{ that is, } a \in A_\vartheta^*. \end{cases}$$

First, consider $\vartheta \in \mathbb{R}$ such that $a \in A'_\vartheta \setminus A_\vartheta^*$, that is, $\ell_w(\vartheta) = \ell_v(\vartheta) + \tau_a$. Due to the continuity of ℓ , the set $\{\theta \in \mathbb{R} \mid a \in A'_\theta \setminus A_\theta^*\}$ consists of closed intervals. Consequently, for almost every considered ϑ it holds $\varrho_w(\vartheta) = \varrho_v(\vartheta)$. Therefore, $f_a^-(\ell_v(\vartheta) + \tau_a) = f_a^-(\ell_w(\vartheta)) = f_a^+(\ell_v(\vartheta))$ and $f_a^-(\ell_w(\vartheta)) \leq \nu_a$. Hence, the equation holds for almost all $\vartheta \in \mathbb{R}$ such that $\varrho_v(\vartheta) \neq 0$ and $a \in A'_\vartheta \setminus A_\vartheta^*$.

If $a \in A \setminus A'_\vartheta$, define $\vartheta := \sup\{\theta \leq \vartheta \mid a \in A'_\theta\} \in \mathbb{R} \cup \{-\infty\}$ and $\bar{\vartheta} := \inf\{\theta \geq \vartheta \mid a \in A'_\theta\} \in \mathbb{R} \cup \{+\infty\}$. Continuity of ℓ yields $a \in A \setminus A'_\theta$ for all $\theta \in (\vartheta, \bar{\vartheta})$ as well as $\ell_w(\vartheta) = \ell_v(\vartheta) + \tau_a$ and $\ell_w(\bar{\vartheta}) = \ell_v(\bar{\vartheta}) + \tau_a$. Therefore, monotonicity and continuity of ℓ_v and ℓ_w imply that f_a^+ and f_a^- vanish almost everywhere on the intervals $(\ell_v(\vartheta), \ell_v(\bar{\vartheta}))$ and $(\ell_w(\vartheta), \ell_w(\bar{\vartheta})) = (\ell_v(\vartheta) + \tau_a, \ell_v(\bar{\vartheta}) + \tau_a)$, respectively. This proves that the equation is fulfilled for almost every $\vartheta \in \mathbb{R}$ such that $a \in A \setminus A'_\vartheta$.

A similar argument shows that $f_a^-(\ell_v(\vartheta) + \tau_a) = \nu_a$ for almost every ϑ with $a \in A_\vartheta^*$. In conclusion, $z_a(\vartheta) = F_a^+(\vartheta) - F_a^-(\vartheta + \tau_a)$ holds on \mathbb{R} .

The earliest times function. Now, it is immediate to see that ℓ indeed is the earliest times function of the feasible flow over time (f^+, f^-) . For $a = (v, w) \in A$ and $\vartheta \in \mathbb{R}$, the above characterization of z_a in terms of ℓ immediately implies $\ell_w(\vartheta) \leq \ell_v(\vartheta) + 1/\nu_a \cdot z_a(\ell_v(\vartheta)) + \tau_a$ with equality if and only if $a \in A'_\vartheta$. This proves that ℓ satisfies the Bellman equations (BE).

The equilibrium condition. To complete the proof, we only have to show that (f^+, f^-) is a dynamic equilibrium. For $a = (v, w) \in A$, ℓ_v and ℓ_w are monotonically nondecreasing. Hence, for every $\vartheta \in \mathbb{R}$, two changes of variables yield

$$\begin{aligned} F_a^+(\ell_v(\vartheta)) &= \int_0^\vartheta f_a^+(\ell_v(\theta)) \cdot \varrho_v(\theta) \, d\theta \\ &= \int_0^\vartheta \chi_a(\theta) \, d\theta \\ &= \int_0^\vartheta f_a^-(\ell_w(\theta)) \cdot \varrho_w(\theta) \, d\theta = F_a^-(\ell_w(\vartheta)). \end{aligned}$$

Lemma 4.2 gives that (f^+, f^-) is a dynamic equilibrium. ■

4.4.2 Dynamic Equilibria for Right-Monotone Inflow

The above theorem suggests to construct a dynamic equilibrium by integrating over thin flows with resetting. The method of Koch and Skutella (2011) does this by implicitly assuming that ℓ is right-linear. This is feasible for piecewise constant inflow rates, as there always is a dynamic equilibrium with that property. For these dynamic equilibria, the functions $\vartheta \mapsto \varrho_v^{G'_\vartheta}(\nu_0(\vartheta))$ are right-constant. We want to consider a more general class of inflow rates.

Definition 4.17 (Monotone functions). We call a function $g \in L^1_{loc}(\mathbb{R})$ **monotonically nondecreasing (nonincreasing)** if there exists a set $N \subseteq \mathbb{R}$ of measure zero such that $g(\xi) \leq g(\hat{\xi})$ ($g(\xi) \geq g(\hat{\xi})$) for all $\xi, \hat{\xi} \in \mathbb{R} \setminus N$ with $\xi \leq \hat{\xi}$.

We call a function g **monotone** if it is monotonically nondecreasing or monotonically nonincreasing. Further, g is **right-monotone (left-monotone)** if for every $\xi \in \mathbb{R}$ there is $\varepsilon > 0$ such that g is monotone on $[\xi, \xi + \varepsilon]$ ($[\xi - \varepsilon, \xi]$).

For a right-monotone inflow rate ν_0 , the map $\vartheta \mapsto \varrho_v^{G'_\vartheta}(\nu_0(\vartheta))$ cannot be expected to be right-constant. Due to the piecewise linear dependency of the thin flows with resetting on the flow value, however, this map is right-monotone. This still allows to use the same method for constructing a dynamic equilibrium as follows.

Theorem 4.18 (α -extension of dynamic equilibria). Let the inflow rate ν_0 be right-monotone. For $\vartheta \geq 0$, let $\ell: (-\infty, \vartheta] \rightarrow \mathbb{R}_{\geq 0}^V$ fulfill the differential equation (DE) on $(-\infty, \vartheta]$. Then there is $\alpha > 0$ such that ℓ can be extended to fulfill it on $(-\infty, \vartheta + \alpha]$.

For the proof of this theorem, we make use of some basic properties of right-monotone functions which we show first. Lemma 4.19 regards the composition of left-/right-monotone functions. Lemma 4.20 relates right-monotone functions to their primitives.

Lemma 4.19 (Composition of left-/right-monotone functions). *Let $h \in L^1_{loc}(\mathbb{R})$ be right-monotone and $g \in L^1_{loc}(\mathbb{R})$ be left- and right-monotone such that their composition $g \circ h$ is well-defined. If h is locally bounded or there exists $v \geq 0$ such that g is monotone on the unbounded intervals $(-\infty, -v)$ and $(v, +\infty)$, then also $g \circ h$ is right-monotone.*

Proof. Let $\xi \in \mathbb{R}$. Assume h is monotonically nondecreasing on $(\xi, \xi + \varepsilon)$ for some $\varepsilon > 0$. Set v to be the essential infimum of h on $(\xi, \xi + \varepsilon)$, that is,

$$v := \text{ess inf } h|_{(\xi, \xi + \varepsilon)} = \sup\{v \in \mathbb{R} \mid h \geq v \text{ a.e. on } (\xi, \xi + \varepsilon)\}.$$

If $v > -\infty$, there is $\bar{v} > v$ such that g is monotone on (v, \bar{v}) as g is right-monotone. Otherwise, h is not locally bounded and the assumption gives $\bar{v} \in \mathbb{R}$ such that g is monotone on $(v, \bar{v}) = (-\infty, \bar{v})$. Choose $0 < \hat{\varepsilon} \leq \varepsilon$ small enough such that $h \leq \bar{v}$ almost everywhere on $(\xi, \xi + \hat{\varepsilon})$. Then $g \circ h$ is monotone on $(\xi, \xi + \hat{\varepsilon})$.

The case that h is monotonically nonincreasing on $(\xi, \xi + \varepsilon)$ works similarly. It needs the left-monotonicity instead of right-monotonicity of g . ■

Lemma 4.20. *Let $g \in L^1_{loc}(\mathbb{R})$ be a right-monotone function such that $G: \mathbb{R} \rightarrow \mathbb{R}, \xi \mapsto \int_0^\xi g(\hat{\xi}) d\hat{\xi}$ satisfies $\inf\{\xi > 0 \mid G(\xi) > 0\} = 0$. Then, there is $\varepsilon > 0$ such that $g > 0$ almost everywhere on $(0, \varepsilon)$ and $G > 0$ on $(0, \varepsilon)$.*

Proof. Let $\varepsilon > 0$ such that g is monotone on $(0, \varepsilon)$, and set

$$\delta := \sup\{\hat{\delta} \in [0, \varepsilon) \mid g \leq 0 \text{ a.e. on } [0, \hat{\delta}]\}.$$

$\delta > 0$ would imply $G \leq 0$ on $[0, \delta]$ which contradicts $\inf\{\xi > 0 \mid G(\xi) > 0\} = 0$. Hence, $\delta = 0$. If g is monotonically nondecreasing on $(0, \varepsilon)$, then $g > 0$ almost everywhere on $(0, \varepsilon)$ and $G > 0$ on $(0, \varepsilon)$ follow. If g is monotonically nonincreasing on $(0, \varepsilon)$, then $\delta = 0$ implies $\text{ess sup } g|_{(0, \varepsilon)} > 0$. Thus, there is $0 < \hat{\varepsilon} \leq \varepsilon$ such that $g > 0$ almost everywhere on $(0, \hat{\varepsilon})$ and $G > 0$ on $(0, \hat{\varepsilon})$. ■

Proof of Theorem 4.18. We extend ℓ by assuming that G'_ϑ is constant to the right, that is, $G'_{\vartheta+\varepsilon} = G'_\vartheta$ for small $\varepsilon > 0$. This assumption is not necessarily true. We will see, however, that the extension we get in this way fulfills (DE).

The extension. For that purpose, define for $\varepsilon > 0$

$$\ell(\vartheta + \varepsilon) := \ell(\vartheta) + \int_\vartheta^{\vartheta+\varepsilon} \varrho^{G'_\vartheta}(\nu_0(\theta)) d\theta,$$

where integration is applied element-wise. First of all, we show that this definition is sound. For component $v \in V$, the integrand $\varrho_v^{G'_\vartheta} \circ \nu_0$ is indeed locally integrable since

for every compact set $K \subseteq \mathbb{R}$, Proposition 4.14 yields

$$\begin{aligned} \int_K |\varrho_v^{G'_\vartheta}(\nu_0(\theta))| d\theta &= \int_{K \cap \nu_0^{-1}([0,1])} \varrho_v^{G'_\vartheta}(\nu_0(\theta)) d\theta + \int_{K \setminus \nu_0^{-1}([0,1])} \frac{\varrho_v^{G'_\vartheta}(\nu_0(\theta))}{\nu_0(\theta)} \cdot \nu_0(\theta) d\theta \\ &\leq \varrho_v^{G'_\vartheta}(1) \cdot \int_K \max\{1, \nu_0(\theta)\} d\theta < +\infty. \end{aligned}$$

The feasibility. This extension defines $G'_{\vartheta+\varepsilon} = (V, A'_{\vartheta+\varepsilon}, A^*_{\vartheta+\varepsilon})$ for $\varepsilon > 0$. To show that ℓ fulfills (DE) on a strictly larger interval than $(-\infty, \vartheta]$, it is sufficient to prove $\varrho^{G'_{\vartheta+\varepsilon}} \circ \nu_0 \equiv \varrho^{G'_\vartheta} \circ \nu_0$ for all small enough $\varepsilon > 0$. For that purpose, define the following limits of the sets of active and resetting arcs,

$$\begin{aligned} A' &:= \liminf_{\varepsilon \rightarrow 0^+} A'_{\vartheta+\varepsilon} = \{a = (v, w) \in A \mid \exists \delta > 0: \ell_w \geq \ell_v + \tau_a \text{ on } (\vartheta, \vartheta + \delta)\} \text{ and} \\ A^* &:= \liminf_{\varepsilon \rightarrow 0^+} A^*_{\vartheta+\varepsilon} = \{a = (v, w) \in A \mid \exists \delta > 0: \ell_w > \ell_v + \tau_a \text{ on } (\vartheta, \vartheta + \delta)\}. \end{aligned}$$

We show that $G'_{\vartheta+\varepsilon} = (V, A', A^*)$ for small $\varepsilon > 0$. Due to the continuity of ℓ , we know that the inclusions $A' \subseteq A'_{\vartheta+\varepsilon} \subseteq A'_\vartheta$ and $A^* \subseteq A^*_{\vartheta+\varepsilon} = A^*_\vartheta$ hold for $\varepsilon > 0$ small enough.

Since ν_0 is right-monotone and $\varrho^{G'_\vartheta}$ is piecewise linear (with finitely many breakpoints), Lemma 4.19 implies that $\varrho_w^{G'_\vartheta} \circ \nu_0 - \varrho_v^{G'_\vartheta} \circ \nu_0$ is right-monotone for all pairs $v, w \in V$.

Active arcs. For arcs $a = (v, w) \in A'_\vartheta \setminus A'$, there is a sequence $(\varepsilon_k)_{k \in \mathbb{N}} \subseteq \mathbb{R}_{>0}$ such that $\lim_{k \rightarrow +\infty} \varepsilon_k = 0$ and $a \in A'_{\vartheta+\varepsilon_k} \setminus A'_{\vartheta+\varepsilon_k}$ for all $k \in \mathbb{N}$. As ℓ is continuous and $a \in A'_\vartheta$, a limit argument yields $\ell_w(\vartheta) = \ell_v(\vartheta) + \tau_a$. Further, for all $k \in \mathbb{N}$, we get

$$\int_{\vartheta}^{\vartheta+\varepsilon_k} \varrho_w^{G'_\vartheta}(\nu_0(\theta)) - \varrho_v^{G'_\vartheta}(\nu_0(\theta)) d\theta = (\ell_w(\vartheta+\varepsilon_k) - \ell_v(\vartheta+\varepsilon_k) - \tau_a) - (\ell_w(\vartheta) - \ell_v(\vartheta) - \tau_a) < 0.$$

Lemma 4.20 implies $\varrho_w^{G'_\vartheta} \circ \nu_0 < \varrho_v^{G'_\vartheta} \circ \nu_0$ and, hence, $\chi_a^{G'_\vartheta} \circ \nu_0 = 0$ almost everywhere on $(\vartheta, \vartheta + \beta_a)$ for some $\beta_a > 0$. Further, it implies $a \notin A'_{\vartheta+\varepsilon}$ for small $\varepsilon > 0$. $A' = A'_{\vartheta+\varepsilon}$ follows for small $\varepsilon > 0$.

Resetting arcs. For $a = (v, w) \in A^* \setminus A^*_\vartheta$, there exists a sequence $(\varepsilon_k)_{k \in \mathbb{N}} \subseteq \mathbb{R}_{>0}$ such that $\lim_{k \rightarrow +\infty} \varepsilon_k = 0$ and $a \in A^*_{\vartheta+\varepsilon_k} \setminus A^*_\vartheta$ for all $k \in \mathbb{N}$. Continuity of ℓ yields $\ell_w(\vartheta) = \ell_v(\vartheta) + \tau_a$. Hence, for $k \in \mathbb{N}$

$$\int_{\vartheta}^{\vartheta+\varepsilon_k} \varrho_w^{G'_\vartheta}(\nu_0(\theta)) - \varrho_v^{G'_\vartheta}(\nu_0(\theta)) d\theta = (\ell_w(\vartheta+\varepsilon_k) - \ell_v(\vartheta+\varepsilon_k) - \tau_a) - (\ell_w(\vartheta) - \ell_v(\vartheta) - \tau_a) > 0.$$

Lemma 4.20 implies $\varrho_w^{G'_\vartheta} \circ \nu_0 > \varrho_v^{G'_\vartheta} \circ \nu_0$ and, hence, $\chi_a^{G'_\vartheta} \circ \nu_0 = \nu_a \cdot \varrho_w^{G'_\vartheta} \circ \nu_0$ almost everywhere on $(\vartheta, \vartheta + \beta_a)$ for some $\beta_a > 0$.

In total, there is an $\alpha > 0$ such that $G'_{\vartheta+\varepsilon} = (V, A', A^*)$ for all $0 < \varepsilon < \alpha$. More importantly, $\varrho^{G'_\vartheta}(\nu_0(\vartheta + \varepsilon))$ are the corresponding labels of a normalized thin flow with resetting of value $\nu_0(\vartheta + \varepsilon)$ in $G'_{\vartheta+\varepsilon}$. Corollary 4.12 yields $\varrho^{G'_\vartheta} \circ \nu_0 \equiv \varrho^{G'_{\vartheta+\varepsilon}} \circ \nu_0$ on $(\vartheta, \vartheta + \alpha)$. \blacksquare

Just like in the case of constant inflow rate, the α -extension can be applied iteratively to construct a dynamic equilibrium. Each such extension that is maximal with respect to α is called a **phase** in the evolution of the dynamic equilibrium. Note that [Theorem 4.18](#) does not state anything about the length of a phase. It is still an open question whether for constant inflow rate the extensions can converge to a finite domain, see Cominetti, Correa, and Olver [2017](#). If that is the case, the dynamic equilibrium cannot be computed this way in a finite number of steps. In theory, we can take the limit of such a converging sequence of α -extensions and repeat. As done by Cominetti, Correa, and Olver ([2017](#)) and Graf and Harks ([2019](#)), the Kuratowski-Zorn lemma ([1922](#), [1935](#)) can be applied to get a dynamic equilibrium on \mathbb{R} .

Theorem 4.21 (Existence of dynamic equilibria). *For every nonnegative, right-monotone, locally integrable inflow rate, there exists a dynamic equilibrium.*

Proof. Let \mathcal{L} be the set of functions $\ell: (-\infty, \vartheta] \rightarrow \mathbb{R}_{\geq 0}^V$ which fulfill (DE) on $(-\infty, \vartheta)$ for some $\vartheta \in \mathbb{R}_{\geq 0} \cup \{+\infty\}$. Define the partial order \preceq on \mathcal{L} by setting $\ell \preceq \hat{\ell}$ if $\text{dom}(\ell) \subseteq \text{dom}(\hat{\ell})$ and $\ell = \hat{\ell}|_{\text{dom}(\ell)}$. Here, $\text{dom}(\cdot)$ denotes the domain of a function, and $\cdot|_X$ denotes the restriction of a function to a subset X of its domain.

Let $(\ell^{(k)})_{k \in K}$ be a chain in (\mathcal{L}, \preceq) indexed by some set K with domains $\text{dom}(\ell^{(k)}) = (-\infty, \vartheta_k]$. Set $\vartheta := \sup_{k \in K} \vartheta_k$ and define the function

$$\ell: (-\infty, \vartheta) \rightarrow \mathbb{R}_{\geq 0}^V, \theta \mapsto \sup_{k \in K: \theta \leq \vartheta_k} \ell^{(k)}(\theta).$$

Note that $\ell|_{(-\infty, \vartheta_k)} \equiv \ell^{(k)}$ for all $k \in K$. We would like to continuously extend ℓ to $\hat{\ell}$ onto $(-\infty, \vartheta]$. Therefore, set $\hat{\ell}(\vartheta) := \lim_{k \rightarrow +\infty} \ell(\vartheta_k)$. This limit exists as ℓ is monotonically nondecreasing. It is not clear, however, that it is finite. Since ℓ fulfills (DE) on $(-\infty, \vartheta)$, applying [Proposition 4.14](#) yields

$$\begin{aligned} \hat{\ell}_v(\vartheta) &= \lim_{k \rightarrow +\infty} \int_0^{\vartheta_k} \varrho_v^{G'_\theta}(\nu_0(\theta)) \, d\theta \\ &= \lim_{k \rightarrow +\infty} \int_{[0, \vartheta_k] \cap \nu_0^{-1}([0, 1])} \varrho_v^{G'_\theta}(\nu_0(\theta)) \, d\theta + \int_{[0, \vartheta_k] \setminus \nu_0^{-1}([0, 1])} \frac{\varrho_v^{G'_\theta}(\nu_0(\theta))}{\nu_0(\theta)} \cdot \nu_0(\theta) \, d\theta \\ &\leq \max_{\substack{G'=(V, A', A^*) \\ \text{shortest path graph} \\ \text{with resetting}}} \varrho_v^{G'}(1) \cdot \int_0^\vartheta \max\{1, \nu_0(\theta)\} \, d\theta < +\infty. \end{aligned}$$

Hence, $\hat{\ell} \in \mathcal{L}$ is well-defined and $\ell^{(k)} \preceq \hat{\ell}$ for all $k \in K$. The Kuratowski-Zorn lemma yields a maximal element $\ell \in \mathcal{L}$. [Theorem 4.18](#) shows that the domain of ℓ has to be \mathbb{R} . By [Theorem 4.16](#), ℓ is the earliest times function of a dynamic equilibrium. \blacksquare

Theorem 4.22 (Uniqueness of right-monotone dynamic equilibria). *Let the inflow rate ν_0 be right-monotone. If there are two right-monotone dynamic equilibria, then their earliest times functions agree.*

Proof. Let $\Theta \subseteq \mathbb{R}$ be the set on which the earliest times functions of every right-monotone equilibrium with inflow rate ν_0 agree. By definition $\mathbb{R}_{\leq 0} \subseteq \Theta$. Since earliest times functions are continuous, Θ is a closed set. We need to show that $\Theta = \mathbb{R}$.

Let (f^+, f^-) be a dynamic equilibrium for inflow rate ν_0 such that f_a^+ is right-monotone for all $a = (v, w) \in A$. Let ℓ be the corresponding earliest times function. For almost every $\vartheta \in \mathbb{R}$, z_a is increasing at ϑ if and only if $f_a^+(\vartheta) > \nu_a$. Therefore, z_a is right-monotone as well. (Looking into the proof of) [Lemma 4.19](#) yields right-monotonicity of $f_a^+ \circ \ell_v$ and $z_a \circ \ell_v$ since ℓ_v is continuous, hence, locally bounded, and nondecreasing.

Let $\vartheta \in \Theta$. We will see that the support of $(f_a^+(\ell_v(\vartheta + \varepsilon)))_{a=(v,w) \in A}$ and $A_{\vartheta+\varepsilon}^*$ are respectively the same for almost every small enough $\varepsilon > 0$. Define the set

$$A' := \{a = (v, w) \in A \mid \exists \delta > 0: f_a^+ \circ \ell_v > 0 \text{ a.e. on } (\vartheta, \vartheta + \delta)\}.$$

For $a \in A \setminus A'$, the right-monotonicity of $f_a^+ \circ \ell_v$ implies $f_a^+(\ell_v(\vartheta + \varepsilon)) = 0$ for almost every small enough $\varepsilon > 0$. Thus, the support of $(f_a^+(\ell_v(\vartheta + \varepsilon)))_{a=(v,w) \in A}$ is exactly A' . Similarly, define the set

$$A^* := \liminf_{\varepsilon \rightarrow 0^+} A_{\vartheta+\varepsilon}^* = \{a = (v, w) \in A \mid \exists \delta > 0: z_a \circ \ell_v > 0 \text{ on } (\vartheta, \vartheta + \delta)\}.$$

As $z_a \circ \ell_v$ is right-monotone, $A^* = A_{\vartheta+\varepsilon}^*$ for small enough $\varepsilon > 0$. Thus for almost every small enough $\varepsilon > 0$, $(f_a^+(\ell_v(\vartheta + \varepsilon)) \cdot \frac{d\ell_v}{d\vartheta}(\vartheta + \varepsilon))_{a=(v,w) \in A}$ is a normalized thin flow of value $\nu_0(\vartheta + \varepsilon)$ in $G' := (V, A', A^*)$ with corresponding labels $\frac{d\ell}{d\vartheta^+}(\vartheta + \varepsilon)$. Hence, [Theorem 4.16](#) implies

$$\ell(\vartheta + \varepsilon) = \ell(\vartheta) + \int_{\vartheta}^{\vartheta+\varepsilon} \varrho^{G'}(\nu_0(\theta)) d\theta.$$

Since (f^+, f^-) was chosen arbitrarily, $[\vartheta, \vartheta + \varepsilon] \subseteq \Theta$. As a consequence, $\Theta = \mathbb{R}$. ■

The following example shows that the α -extension may fail if ν_0 is not right-monotone.

Example 4.23 (α -extension for non-right-monotone inflow rate). We consider the graph $G = (V, A)$ which consists only of the source s , the sink t , and two parallel arcs a, b from s to t with capacity $\nu_a = \nu_b = 1$ and transit time $\tau_a = 0$ and $\tau_b = 1$, see [Figure 4.6a](#). As mentioned before the network can be transformed to an equivalent instance without multi-arcs. The inflow rate is depicted in [Figure 4.6b](#) and given by the function

$$\nu_0: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}, \quad \vartheta \mapsto \begin{cases} 0 & \text{for } \vartheta \in (-\infty, 0) \\ 2 & \text{for } \vartheta \in [0, 1] \\ 0 & \text{for } \vartheta \in [1 + 2^{-k-1}, 1 + 2^{-k}), k \in 2\mathbb{N} + 1. \\ 2 & \text{for } \vartheta \in [1 + 2^{-k-1}, 1 + 2^{-k}), k \in 2\mathbb{N} \\ 2 & \text{for } \vartheta \in [2, +\infty) \end{cases}$$

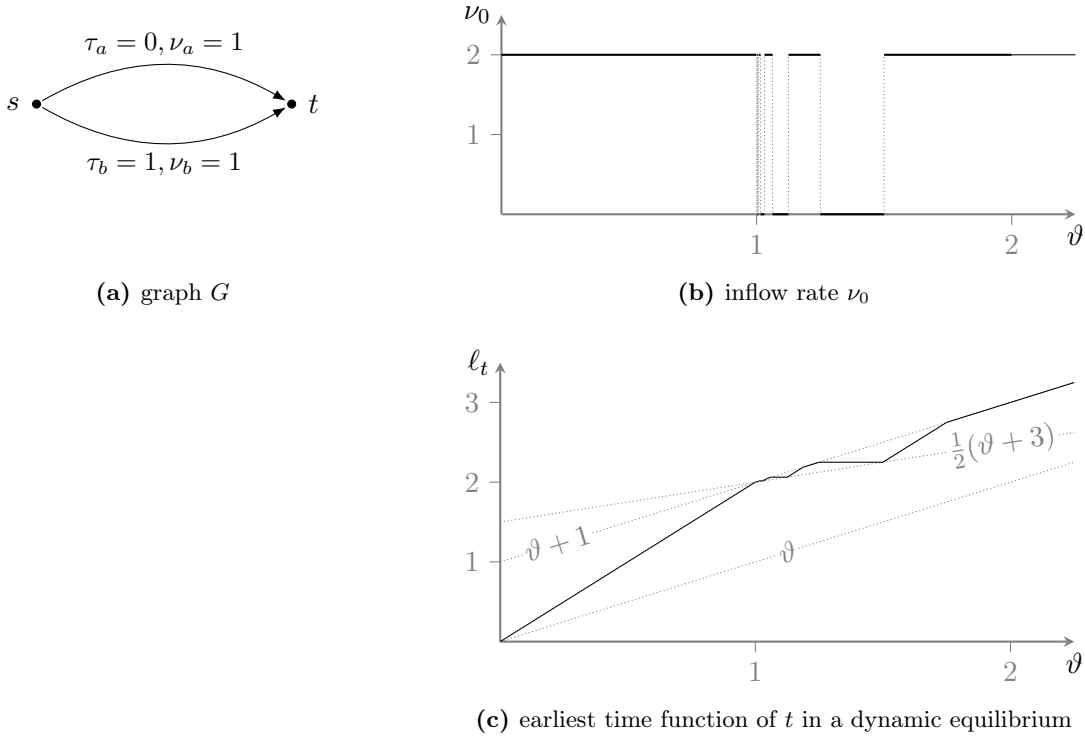


Figure 4.6: Example with nonright-monotone inflow rate for which the α -extension does not work.

Note that $\nu_0 \in L_{\text{loc}}^1(\mathbb{R}_{\geq 0})$ is not right-monotone at $\vartheta = 1$. A dynamic equilibrium is given by the earliest time functions $\ell_s: \mathbb{R} \rightarrow \mathbb{R}, \vartheta \mapsto \vartheta$ and

$$\ell_t: \mathbb{R} \rightarrow \mathbb{R}, \quad \vartheta \mapsto \begin{cases} \vartheta & \text{for } \vartheta \in (-\infty, 0) \\ 2\vartheta & \text{for } \vartheta \in [0, 1] \\ 2 + 2^{-k-1} & \text{for } \vartheta \in [1 + 2^{-k-1}, 1 + 2^{-k}), k \in 2\mathbb{N} - 1 \\ 2\vartheta - 3 \cdot 2^{-k-2} & \text{for } \vartheta \in [1 + 2^{-k-1}, 1 + 3 \cdot 2^{-k-2}), k \in 2\mathbb{N} \\ \vartheta + 1 & \text{for } \vartheta \in [1 + 3 \cdot 2^{-k-2}, 1 + 2^{-k}), k \in 2\mathbb{N} \\ \vartheta + 1 & \text{for } \vartheta \in [2, +\infty) \end{cases}$$

The graph of ℓ_t is shown in [Figure 4.6c](#). A queue grows on a up to time 1, when b becomes active, that is, $A'_1 = \{a, b\}$. After time 1, the following three phases repeat for every $k \in 2\mathbb{N}$. For times $\vartheta \in [1 + 2^{-k-2}, 1 + 2^{-k-1})$, there is no inflow, the set of active arcs is $A'_\vartheta = \{a\}$, and the queue on a shrinks. At time $1 + 2^{-k-1}$ the inflow sets in again. For times $\vartheta \in [1 + 2^{-k-1}, 1 + 3 \cdot 2^{-k-2})$, the set of active arcs is still $A'_\vartheta = \{a\}$, but the queue on a is growing. At time $1 + 3 \cdot 2^{-k-2}$, b gets active. For times $\vartheta \in [1 + 3 \cdot 2^{-k-2}, 1 + 2^{-k})$, the set of active arcs is $A'_\vartheta = \{a, b\}$ and the queue lengths stay constant. At time $1 + 2^{-k}$,

the inflow stops again, and the three phases repeat. In particular, G'_ϑ is not constant on $(1, 1 + \varepsilon)$ for any $\varepsilon > 0$. Applying an α -extension at time 1 yields for small $\varepsilon > 0$

$$\ell_t: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}, \quad 1 + \varepsilon \mapsto \begin{cases} 1 + \frac{1}{3}2^{-k} & \text{for } \varepsilon \in [2^{-k-1}, 2^{-k}), k \in 2\mathbb{N} - 1 \\ 1 + \varepsilon - \frac{1}{3}2^{-k} & \text{for } \varepsilon \in [2^{-k-1}, 2^{-k}), k \in 2\mathbb{N} \end{cases}.$$

Then, $A'_{1+\varepsilon} = \{a\}$ for small $\varepsilon > 0$. But (DE) is not fulfilled at times $1 + \varepsilon$ for all $\varepsilon \in [2^{-k-1}, 2^{-k}), k \in 2\mathbb{N}$.

4.5 Thin Flows with Resetting in Series-Parallel Graphs

The nature of the proof of existence makes the problem of computing a normalized thin flow with resetting part of the complexity class PPAD. It remains open whether this computational problem can be solved in polynomial time for general acyclic arc sets $A^* \subseteq A' \subseteq A$. So far, the complexity is known only for the cases $A^* = A'$ and $A^* = \emptyset$. For $A^* = A'$, the complementarity problem in [Theorem 4.5](#) reduces to a system of linear equations and, hence, can be solved efficiently. For $A^* = \emptyset$, the problem is also solvable in polynomial time as shown by Koch and Skutella ([2011](#)). Instead of restricting the set A^* , we address the computation of normalized thin flows with resetting for graphs (V, A') that are series-parallel.

The following two lemmas show how the function $\varrho^{G'}$ relates to the functions $\varrho^{G'_1}$ and $\varrho^{G'_2}$ for the series composition $G' = G'_1 * G'_2$ as well as the parallel composition $G' = G'_1 \parallel G'_2$. Restricting a thin flow with resetting in G' to G'_1 or G'_2 yields a thin flow with resetting in the respective graph. The other way round, the conditions for a flow to be a thin flow with resetting in G' are roughly a combination of the conditions for flows in G'_1 and G'_2 . The only additional requirement is to have one common label ℓ'_v at every common vertex v of G'_1 and G'_2 . For the series composition this synchronization of labels is straightforward. For the parallel composition, it is achieved by splitting the total amount of flow appropriately between G'_1 or G'_2 .

Lemma 4.24 (Series composition). *Let $G' = (V, A')$ be a two-terminal directed acyclic graph. Further, let $V_1, V_2 \subseteq V$ be vertex sets and set $G'_i := G'[V_i]$ for $i = 1, 2$. If $G' = G'_1 * G'_2$ and n_1, n_2, n are the numbers of breakpoints of $\varrho^{G'_1}, \varrho^{G'_2}, \varrho^{G'}$, respectively, then $n \leq n_1 + n_2$.*

Proof. Let $r \in V$ be such that $V_1 \cap V_2 = \{r\}$. As discussed in the paragraph before [Lemma 4.24](#), the conditions of [Definition 4.4](#) for G'_1 and G'_2 also appear for G' . Applying [Lemma 4.15](#) yields that, for $\nu'_0 > 0$

$$\begin{aligned} \varrho_v^{G'}(\nu'_0) &= \varrho_v^{G'_1}(\nu'_0) && \text{if } v \in V_1 \text{ and} \\ \varrho_v^{G'}(\nu'_0) &= \varrho_r^{G'_1}(\nu'_0) \cdot \varrho_v^{G'_2} \left(\frac{\nu'_0}{\varrho_r^{G'_1}(\nu'_0)} \right) && \text{if } v \in V_2. \end{aligned}$$




Figure 4.7: Series composition $G' = G'_1 * G'_2$ at common vertex r . The breakpoint of $\varrho_t^{G'}$ at $\hat{\nu}_0$ results from the breakpoint of $\varrho_t^{G'_2}$.

For $v \in V_2$, $\varrho_v^{G'}$ can have a breakpoint at $\nu'_0 > 0$ only if $\varrho_r^{G'_1}$ does, or $\varrho_v^{G'_2}$ has a breakpoint at $\frac{\nu'_0}{\varrho_r^{G'_1}(\nu'_0)}$ and the function $\frac{\nu'_0}{\varrho_r^{G'_1}(\nu'_0)}$ is not constant around ν'_0 . Lemma 4.15 shows that $\frac{\nu'_0}{\varrho_r^{G'_1}(\nu'_0)}$ is monotone in ν'_0 . Therefore, $n \leq n_1 + n_2$ follows. See Figure 4.7 for an illustration. ■

Lemma 4.25 (Parallel composition). *Let $G' = (V, A')$ be a two-terminal directed acyclic graph. Further, let $V_1, V_2 \subseteq V$ be vertex sets and set $G'_i := G'[V_i]$ for $i = 1, 2$. If $G' = G'_1 \parallel G'_2$ and n_1, n_2, n are the numbers of breakpoints of $\varrho^{G'_1}, \varrho^{G'_2}, \varrho^{G'}$, respectively, then $n \leq n_1 + n_2 + 1$.*

Proof. Considering Definition 4.4 or the complementarity problem from Theorem 4.5 reveals that for every $\nu'_0 \geq 0$

$$\begin{aligned} \varrho_t^{G'}(\nu'_0) &= \min_{i \in \{1, 2\}} \varrho_t^{G'_i}(\nu'_i) \\ \varrho_t^{G'}(\nu'_0) &= \varrho_t^{G'_i}(\nu'_i) && \text{for } i \in \{1, 2\} \text{ with } \nu'_i > 0 \\ \varrho_v^{G'}(\nu'_0) &= \varrho_v^{G'_i}(\nu'_i) && \text{for } i \in \{1, 2\}, v \in V_i \setminus \{t\} \\ \nu'_1 + \nu'_2 &= \nu'_0 \\ \nu'_1, \nu'_2 &\geq 0. \end{aligned}$$




Figure 4.8: Parallel composition $G' = G'_1 \parallel G'_2$. The breakpoint $\hat{\nu}_0$ of $\varrho_t^{G'}$ is based on the intersection of $\varrho_t^{G'_1}(\nu'_0)$ and $\varrho_t^{G'_2}(\hat{\nu}_0 - \nu'_0)$ at a breakpoint of $\varrho_t^{G'_1}$.

By Corollary 4.12, $\varrho_t^{G'}(\nu'_0)$ is uniquely determined by this nonlinear complementarity problem for given ν'_0 , $\varrho_t^{G'_1}$, and $\varrho_t^{G'_2}$. However, ν'_1 and ν'_2 are generally not uniquely determined. This fact is independent of the equations for all $v \in V \setminus \{t\}$. Hence, we ignore those for the time being and find functions $\nu_1: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and $\nu_2: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ which are piecewise linear and describe a solution $\nu_1(\nu'_0), \nu_2(\nu'_0)$ for every $\nu'_0 \geq 0$. The idea is to define them on a discrete set for which they are unique solutions and interpolate linearly.

For $i \in \{1, 2\}$, define the functions $\underline{\nu}_i: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and $\bar{\nu}_i: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ by

$$\underline{\nu}_i(\rho) := \min \{ \nu'_i \geq 0 \mid \varrho_t^{G'_i}(\nu'_i) \geq \rho \} \quad \text{and} \quad \bar{\nu}_i(\rho) := \inf \{ \nu'_i \geq 0 \mid \varrho_t^{G'_i}(\nu'_i) > \rho \}.$$

Then for $\rho \geq \varrho_t^{G'_i}(0)$, the preimage of ρ under the function $\varrho_t^{G'_i}$ is exactly the interval $[\underline{\nu}_i(\rho), \bar{\nu}_i(\rho)]$. For $\rho < \varrho_t^{G'_i}(0)$, we get $\underline{\nu}_i(\rho) = \bar{\nu}_i(\rho) = 0$. Note that $\bar{\nu}_i(\rho) < \underline{\nu}_i(\hat{\rho})$ for all $\varrho_t^{G'_i}(0) \leq \rho < \hat{\rho}$ due to the monotonicity of $\varrho_t^{G'_i}$.

Without loss of generality, we can assume that $\varrho_t^{G'_1}(0) \leq \varrho_t^{G'_2}(0)$. For $\nu'_0 \geq 0$, there is a unique $\rho \geq \varrho_t^{G'_1}(0)$ such that $\nu'_0 \in [\nu_1(\rho) + \nu_2(\rho), \bar{\nu}_1(\rho) + \bar{\nu}_2(\rho)]$. If $\rho < \varrho_t^{G'_2}(0)$, then $\rho = \varrho_t^{G'_1}(\nu'_0)$. Otherwise, ρ is the value at the intersection of the graphs $\varrho_t^{G'_1}$ and $\varrho_t^{G'_2}(\nu'_0 - \cdot)$, see Figure 4.8. The set of solutions $(\varrho_t^{G'}(\nu'_0), \nu'_1, \nu'_2)$ to the above complementarity problem

can be written as

$$\{(\rho, \nu'_1, \nu'_2) \in \mathbb{R}_{\geq 0} \times [\nu_1(\rho), \bar{\nu}_1(\rho)] \times [\nu_2(\rho), \bar{\nu}_2(\rho)] \mid \nu'_1 + \nu'_2 = \nu'_0, \rho \geq \varrho_t^{G'}(0)\}.$$

Thus, a solution (ρ, ν'_1, ν'_2) is unique if and only if $\nu'_0 = \nu_1(\rho) + \nu_2(\rho)$ or $\nu'_0 = \bar{\nu}_1(\rho) + \bar{\nu}_2(\rho)$.

Let $R := \{\varrho_t^{G'_1}(0), \varrho_t^{G'_2}(0)\} \cup \{\varrho_t^{G'_i}(\nu'_0) \mid i \in \{1, 2\} \text{ and } \nu'_0 \geq 0 \text{ is breakpoint of } \varrho_t^{G'_i}\}$ be the values at breakpoints of $\varrho_t^{G'_1}$ and $\varrho_t^{G'_2}$ (including the border of the domain). By the above, ν_1 and ν_2 are in particular uniquely determined on $N := \{\nu_1(\rho) + \nu_2(\rho) \mid \rho \in R\} \cup \{\bar{\nu}_1(\rho) + \bar{\nu}_2(\rho) : \rho \in R\}$. Hence, for every $i \in \{1, 2\}$ and $\rho \in R$, we set

$$\nu_i(\nu_1(\rho) + \nu_2(\rho)) := \nu_i(\rho) \quad \text{and} \quad \nu_i(\bar{\nu}_1(\rho) + \bar{\nu}_2(\rho)) := \bar{\nu}_i(\rho).$$

These definitions are extended to $\mathbb{R}_{\geq 0}$ by linear interpolation. Then for $i = 1, 2$, the composition $\varrho_t^{G'_i} \circ \nu_i$ is piecewise linear with breakpoints only in N . In particular, all breakpoints of $\varrho_t^{G'}$ lie in N . By construction, $\nu_1(\nu'_0)$ and $\nu_2(\nu'_0)$ define a solution for every $\nu'_0 \in N$. Linearity in-between the breakpoints generalizes this to all $\nu'_0 \in \mathbb{R}_{\geq 0}$.

The above allows to bound the number of breakpoints of $\varrho_t^{G'}$. For $\rho \in R$, the strict inequality $\nu_1(\rho) + \nu_2(\rho) < \bar{\nu}_1(\rho) + \bar{\nu}_2(\rho)$ holds only if there is $i \in \{1, 2\}$ such that $\nu_i(\rho) < \bar{\nu}_i(\rho)$ and, thus, $\nu_i(\rho)$ and $\bar{\nu}_i(\rho)$ are two breakpoints of $\varrho_t^{G'_i}$ with value ρ . This shows the inequality $|N| \leq n_1 + n_2 + 2$, where the constant is accounting for the border of the domain. As $0 \in N$ due to $\varrho_t^{G'_1}(0) \in R$, it follows that $\varrho_t^{G'}$ has at most $n_1 + n_2 + 1$ breakpoints. Repeating the counting more carefully allows to extend this bound to the number of breakpoints of $\varrho^{G'}$. The set N only contains breakpoints which are based on breakpoints of $\varrho_t^{G'_1}$ and $\varrho_t^{G'_2}$. For $i = 1, 2$, any breakpoint of $\varrho^{G'_i}$ at which $\varrho_t^{G'_i}$ is differentiable leads to at most one breakpoint of $\varrho^{G'_i} \circ \nu_i$ additional to those in N . Consequently, $\varrho^{G'}$ does not have more than $n_1 + n_2 + 1$ breakpoints. ■

Understanding both compositions that series-parallel graphs are based on allows us to compute the function $\varrho^{G'}$ for these graphs recursively. Our algorithm not only computes a normalized thin flow with resetting for one single flow value, but for all flow values simultaneously. As seen in [Section 4.4](#), this is necessary in order to compute dynamic equilibria for more general inflow rates than piecewise constant functions.

Theorem 4.26 (Labels in series-parallel graphs). *Let $G' := (V, A')$ be a two-terminal series-parallel directed graph and $A^* \subseteq A'$. Then $\varrho^{G'}$ has at most $2|A'| - |A^*| - |V| + 1$ many breakpoints and can be computed in polynomial time.*

Proof. We prove the claim by a structural induction on two-terminal series-parallel graphs. For the base case, assume $V = \{s, t\}$ and A' contains a single arc $a = (s, t)$. Then, $\varrho_s^{G'} \equiv 1$ and $\varrho_t^{G'} \equiv \varrho^a(1, \cdot)$ as defined in [Definition 4.4](#). Hence, $\varrho^{G'}$ has exactly $1 - |A^*| = 2|A'| - |A^*| - |V| + 1$ many breakpoints. Assume there are $V_1, V_2 \subseteq V$ such that $G' = G'[V_1] * G'[V_2]$. The statement for G' follows immediately from the induction hypothesis and [Lemma 4.24](#) as $|V| = |V_1| + |V_2| - 1$. Now, assume there are $V_1, V_2 \subseteq V$ such that $G' = G'[V_1] \parallel G'[V_2]$. Then, the statement for G' follows immediately from the induction hypothesis and [Lemma 4.25](#) as $|V| = |V_1| + |V_2| - 2$.

Since all involved functions are piecewise linear with few breakpoints, they can be represented efficiently by their linear pieces. Regarding the proofs of [Lemmas 4.24](#) and [4.25](#), it becomes evident that in both cases $\varrho^{G'}$ can be constructed from $\varrho^{G'_1}$ and $\varrho^{G'_2}$ efficiently. To compose G' from single arcs, $|A'| - 1$ compositions are needed. In total, the function $\varrho^{G'}$ can be computed in polynomial time in the size of the input G' , A^* , and $(\nu_a)_{a \in A'}$. ■

If the corresponding labels ℓ' of a normalized thin flow with resetting are known, computing flow values is not hard. The flow value x'_a on $a = (v, w) \in A'$ lies in the interval $[0, \nu_a \ell'_w]$. If $\ell'_v > \ell'_w$ and $a \notin A^*$, then x'_a is zero. If $\ell'_v < \ell'_w$ or $a \in A^*$, then $x'_a = \nu_a \ell'_w$ holds. Finding a flow satisfying these conditions can be done by a simple flow computation.

Corollary 4.27 (Normalized thin flows in series-parallel graphs). *Normalized thin flows with resetting in two-terminal series-parallel graphs can be computed in polynomial time.*

4.6 Closing Remarks

We examine normalized thin flows with resetting as a parametric problem in dependency on the flow value. The results allow us to give a constructive proof for the existence of dynamic equilibria for single-source single-sink networks with right-monotone inflow rate. Further, we obtain a polynomial-time algorithm for computing thin flows with resetting on two-terminal series-parallel networks. The recursive approach that we take for the latter does not seem to generalize to arbitrary networks. A central aspect that we use in this recursion is that every considered subnetwork has a single source and a single sink.

Major open questions on the model, like the price of anarchy and the number of phases, seem to require insights into the relation between the thin flows of subsequent phases. These, in turn, would need a better understanding of the dependency of thin flows with resetting on the sets of active and resetting arcs.

The characterization of normalized thin flows with resetting by a linear complementarity problem in [Theorem 4.5](#) allows us to deduce some basic properties of the functions $\varrho^{G'}$. Beyond this, it opens the way to approach normalized thin flows with resetting through the existing machinery of linear complementarity problems. The properties of the presented linear complementarity problem as established in [Lemma 4.8](#) and in the proof of [Theorem 4.10](#) show that normalized thin flows with resetting can be computed in finitely many steps via Lemke's algorithm. While there is no immediate nontrivial bound on the number of steps in theory, it suggests an efficient method in practice. The complexity of computing normalized thin flows with resetting (or even parametric thin flows) remains open in general. Analyzing the number of iteration of Lemke's algorithm for the proposed linear complementarity problem or similar formulations might shed some light on this question.

Chapter 5

Nash Equilibria in Network Cost-Sharing Games

This chapter treats Nash equilibria in **network cost-sharing games**. These games represent the game theoretic counterpart of buy-at-bulk network design. In network cost-sharing games, a set of players needs to connect their respective source and sink by selecting a path. The use of an edge incurs a cost that is nondecreasing with the number of users. Under fair cost allocation, the cost is fairly split among all users. We restrict to the cases where the total edge cost is **nondecreasing and concave**, and the resulting cost per player is **nonincreasing** in the number of players. We use the term cost-sharing in order to emphasize the latter. Formally, cost-sharing games fall into the class of congestion games. We assume that all players act selfishly and, thus, try to minimize the total cost they spend on connecting their source and sink. As opposed to congestion games with nondecreasing player cost, the sharing of the costs incentivizes coordination between the players up to some point. A defecting player, however, can obliterate the benefits of a group's cooperation. Maybe the most prominent questions in this area of research are the **computational complexity** and the **efficiency** of Nash equilibria. We want to advance the understanding by improving the existing bounds on the price of stability and generalize many existing results to more **general cost functions**.

The questions at hand are closely related to the topic of **minimum concave-cost network flows**. This model differs from the well known minimum-cost flow problem by allowing a separable concave objective function with respect to an edge flow. We highlight this connection in [Section 5.3.1](#). In general, however, we keep the notation of the original domain in order to be more consistent with the existing literature on the topic.

Authorship. The presented results are joint work with Yiannis Giannakopoulos and Clara Waldmann.

Outline. [Section 5.1](#) starts with introducing notation and terminology specific to network cost-sharing games. Further, it relates our contributions to existing literature. In [Section 5.2](#), we establish basic structural results for Nash equilibria and social optima. [Section 5.3](#) deals with the computational complexity of Nash equilibria. On the positive side, we discuss various formulations of the problem of finding a Nash equilibrium in [Section 5.3.1](#). On the negative side, we provide several intractability results in [Sec-](#)

tion 5.3.2. Section 5.4 focuses on the efficiency of Nash equilibria. After analyzing the price of anarchy in Section 5.4.1, we move on to the analysis of the price of stability. Section 5.4.2 provides upper bounds on the price of stability via two different methods. This is augmented by lower bounds for various classes of cost functions in Section 5.4.3. We close this chapter by some final remarks in Section 5.5.

5.1 Introduction

We start by introducing basic terms and notation to define the problem formally. Further, we establish some special classes of instances in terms of the the sources and sinks as well as the cost functions.

Network cost-sharing games. An (undirected) **network cost-sharing game** is defined on an (undirected) graph $G = (V, E)$. There are $n \in \mathbb{N}$ players. Each player $i \in [n]$ wants to connect her source $s_i \in V$ to her sink $t_i \in V$. Therefore, the **strategies** of a player $i \in [n]$ are the sets of edges in E that connect her source s_i and sink t_i , that is,

$$\Sigma_i := \{E' \subseteq E \mid s_i \text{ is connected to } t_i \text{ in } E'\}.$$

We denote the set of all **strategy profiles** by $\Sigma := \times_{i \in [n]} \Sigma_i$. The set of edges that is used by all players in a strategy profile σ is called its **support** and denoted $\text{supp}(\sigma) := \bigcup_{i \in [n]} \sigma_i$. The **congestion** that a strategy profile $\sigma \in \Sigma$ creates on the edges is expressed by the function

$$n_\sigma: E \rightarrow \mathbb{N}, \quad e \mapsto |\{i \in [n] \mid e \in \sigma_i\}|.$$

Depending on the congestion, the players incur costs for the usage of edges. These costs are described by $c \in \mathbb{R}_{\geq 0}^E$ and a positive, nonincreasing **cost function** $f: \mathbb{N} \rightarrow \mathbb{R}_{> 0}$. The requirement of f being nonincreasing distinguishes cost-sharing games from congestion games with increasing cost function. The **total edge cost** that is incurred by $k \in \mathbb{N}$ players using it is given by $c_e k f(k)$ if $k > 0$ and zero otherwise. The total cost is split fairly among all players that use it. This allocation of total cost to users is called the **Shapley cost-sharing mechanism** as the player costs relate to the Shapley value named after Shapley (1951). We define the cost $c_\sigma(e)$ for using an edge $e \in E$ in the strategy profile σ as

$$c_\sigma: E \rightarrow \mathbb{R}_{\geq 0}, \quad e \mapsto c_e f(n_\sigma(e)).$$

For subsets of edges $E' \subseteq E$, we write $c_\sigma(E') := \sum_{e \in E'} c_\sigma(e)$. In a strategy profile $\sigma \in \Sigma$, the **total player cost** of $i \in [n]$ is the sum of all costs she incurs for using edges in σ_i ; that is, the total player cost evaluates to

$$C_i(\sigma) = \sum_{e \in \sigma_i} c_\sigma(e) = \sum_{e \in \sigma_i} c_e f(n_\sigma(e)).$$

We assume that every player strives for minimizing her total player cost.

It is a natural assumption that the total cost of a resource is nondecreasing in the number of users. Hence, we will assume throughout that the total edge cost $c_e k f(k)$ is nondecreasing in k for all $e \in E$. We further assume that the total edge cost exhibits **economies of scale**. Then $k \mapsto c_e k f(k)$ is concave, that is, the marginal total edge cost $c_e(k+1)f(k+1) - c_e k f(k)$ is nonincreasing in $k \in \mathbb{N}$.

The following variant of an edge's cost is very helpful when it comes to examining unilateral deviations. We define $c'_\sigma(e)$ for $e \in E$ as the cost a player $i \in [n]$ with $e \notin \sigma_i$ would incur for using e (assuming no other player changes their strategy). It evaluates to $c'_\sigma(e) = c_e f(n_\sigma(e) + 1)$. Note that $c'_\sigma(e) \leq c_\sigma(e)$ holds for all $e \in E$ as we assume that f is nonincreasing. Again, we write $c'_\sigma(E') := \sum_{e \in E'} c'_\sigma(e)$ for a set of edges $E' \subseteq E$. If $\hat{\sigma} \in \Sigma$ is the result of a player $i \in [n]$ unilaterally deviating from a strategy profile $\sigma \in \Sigma$, then $n_{\hat{\sigma}}(e) = n_\sigma(e)$ for all $e \in \hat{\sigma}_i \cap \sigma_i$ whereas $n_{\hat{\sigma}}(e) = n_\sigma(e) + 1$ for all $e \in \hat{\sigma}_i \setminus \sigma_i$. Thus, the change of the total cost of player i can be expressed as

$$C_i(\hat{\sigma}) - C_i(\sigma) = c_{\hat{\sigma}}(\hat{\sigma}_i) - c_\sigma(\sigma_i) = c'_\sigma(\hat{\sigma}_i \setminus \sigma_i) - c_\sigma(\sigma_i \setminus \hat{\sigma}_i).$$

The **social cost** $C(\sigma)$ of a strategy profile $\sigma \in \Sigma$ is defined as the sum of the total player costs and evaluates to

$$C(\sigma) = \sum_{i \in [n]} C_i(\sigma_i) = \sum_{e \in \text{supp}(\sigma)} c_e n_\sigma(e) f(n_\sigma(e)).$$

As mentioned before, network cost-sharing games belong to the class of congestion games. For these, Rosenthal (1973) shows the existence of Nash equilibria by an elegant potential function argument. Monderer and Shapley (1996) characterize congestion games as all games which have such a potential. For the class of network cost-sharing games, the potential function is defined as follows. Set $F: \mathbb{N} \rightarrow \mathbb{R}, k \mapsto \sum_{l=1}^k f(l)$. Note that $F(0) = 0$. Then the **potential function** is given by

$$\Phi: \Sigma \rightarrow \mathbb{R}, \quad \sigma \mapsto \sum_{e \in E} c_e F(n_\sigma(e)).$$

Classes of network cost-sharing games. Various specializations and variants of network cost-sharing games appear in the literature. The most important ones are determined by the network topology and the type of considered cost functions. The above definition of network cost-sharing games naturally translates to **directed** networks. There is a large difference in the behavior of the directed and undirected versions. If not stated explicitly, we refer to the **undirected** problem throughout the chapter.

With respect to the structure of the sources and sinks, there are three classes: the general case with arbitrary source-sink pairs, multicast, and broadcast. A **multicast game** is a network cost-sharing game in which all players share a common source vertex. We refer to this common source as the **root**. In these games, we identify the players




Figure 5.1: The relevant cost functions under constant total edge cost (yellow), polynomial total edge cost with $\alpha = 0.6$ (violet), affine total edge cost with $s = 0.4$ (green), and linear total edge cost (red).

with their sink. At first glance, this does not allow to represent multicast games with multiple players having the same sink. This can be remedied by modifying the graph and adding additional sinks which are connected by zero-cost edges to the original sinks. A **broadcast game** is a network cost-sharing game with every vertex of the network being associated with a player who wants to connect to a common source. In other words, these are multicast games with the set of players being V . As the empty strategy is dominant for the player associated with the root, she is often excluded from the set of players. Both definitions are equivalent.

When it comes to the cost functions, there is the basic distinction between the uniform and nonuniform settings. If not stated otherwise, we consider the **uniform** case, that is, the cost of every edge varies only in c_e but not in f . If additionally f is allowed to depend on the edge, the game is called **nonuniform** (and $c_e \equiv 1$ can be assumed).

In the following sections, we will particularly discuss four specific choices for the cost function f . Their cost functions are depicted in Figure 5.1. The most prominent choice is $f(k) = 1/k$. We refer to this as **constant total edge cost**, because $c_e k f(k) = c_e$ for $k > 0$. This choice is also called step function pricing in buy-at-bulk network design. It represents the extreme case of f decreasing as fast as possible under the assumption that $k \mapsto kf(k)$ is nondecreasing.

The opposite extreme is given by **linear total edge cost** $k \mapsto kc_e$. Then f does not decrease at all, that is, $f \equiv 1$. In this case, there is no effect of sharing. Therefore, any shortest s_i - t_i path is a dominant strategy for player $i \in [n]$.

We consider two natural interpolations between these two extreme cases. The first is a linear interpolation and given by the choice $f(k) = s + (1-s)/k$ for the parameter $0 \leq s \leq 1$. This corresponds to **affine total edge cost** $k \mapsto c_e + (k-1)sc_e$, which is also called offset linear pricing. For the usage of an edge a startup cost of c_e is payed and any additional use has a unit cost of sc_e . Secondly, we discuss a geometric interpolation given by the class of polynomially decreasing functions $f(k) = k^{\alpha-1}$ for the parameter $0 \leq \alpha \leq 1$. In this case, we obtain **polynomial total edge cost** $k \mapsto c_e k^\alpha$.

5.1.1 Related Literature

We give an overview of related work on network cost-sharing games and connected topics. Primarily, we are interested in understanding Nash equilibria and how their social compares to a social optimum. Finding a social optimum is known as the buy-at-bulk network design problem. Any Nash equilibrium constitutes a feasible solution to the network design problem. Hence, hardness results for the approximation of network design immediately translate to hardness of computing Nash equilibria of the respective quality. The other way round, being able to compute Nash equilibria of small social cost efficiently would imply an approximation algorithm for network design. This suggests that the state of the art for the approximation of network design gives a natural bound for the computation of good Nash equilibria.

The buy-at-bulk network design problem can be viewed as a minimum concave-cost flow problem. Research on the latter focuses on directed networks. Still there are some connections which are relevant for undirected networks. There is a significant difference in the complexity of the discussed problems in directed and undirected networks. The directed version is generally harder. This is the reason why the work on buy-at-bulk network design focuses on the undirected case.

Buy-at-bulk network design. There is a large body of literature on the network design problem. In terms of our problem, network design asks for a social optimum. As solving the problem exactly is hard in most cases, research has focused on approximation algorithms. We only give a selection of results, a survey on approximation techniques and further results is given by Gupta and Könemann (2011).

The work of Awerbuch and Azar (1997), Bartal (1998), and Fakcharoenphol, Rao, and Talwar (2003, 2004) yields an $\mathcal{O}(\log n)$ -approximation algorithm for (uniform, undirected) buy-at-bulk network design. The same guarantee can be achieved when nonuniform costs are allowed but a common source of all players is assumed as shown by Chekuri, Khanna, and Naor (2001) and Meyerson, Munagala, and Plotkin (2000, 2008). For multicast network design with uniform cost, there is a constant factor approximation due to Guha, Meyerson, and Munagala (2001), Gupta, Kumar, Pál, and Roughgarden (2007), Gupta, Kumar, and Roughgarden (2003), and Talwar (2002).

On the negative side, Andrews (2004) shows that there is no $\mathcal{O}((\log n)^{1/2-\varepsilon})$ -approximation algorithm for nonuniform buy-at-bulk network design for any $\varepsilon > 0$, unless all problems in NP of size k can be solved in time $\mathcal{O}(k^{\text{polylog}(k)})$ using randomization. The hard instances use (nonuniform) affine total edge costs. Further, there is no $\mathcal{O}((\log n)^{1/4-\varepsilon})$ -approximation algorithm for uniform buy-at-bulk network design for any $\varepsilon > 0$ under the same assumption. The involved construction uses an edge cost function of the form $k \mapsto \min\{sk, S + k\}$ where $s, S \in \mathbb{R}_{\geq 0}$. For the multicast setting, Chuzhoy, Gupta, Naor, and Sinha (2005, 2008) show that there cannot be an $\mathcal{O}(\log \log n)$ -approximation algorithm, unless all problems in NP of size k are solvable in time $k^{\mathcal{O}(\log \log \log k)}$.

Minimum concave-cost flow. The minimum concave-cost flow problem asks for a flow in a network with supply and demand of minimal total cost with respect to a separable concave objective function. It was formally introduced in the literature by Zangwill (1968). It is essentially equivalent to the buy-at-bulk network design problem. The complexity of the minimum concave-cost flow problem depends on the number of sources and the number of arcs with nonlinear cost function. It is well known that capacities in network flows can be modeled by introducing additional sources and sinks. The following concerns uncapacitated networks after such a transformation. Erickson, Monma, and Veinott (1987) provide an algorithm based on dynamic programming, which they call send-and-split. Its running time is polynomial in the size of the network, but exponential in the number of sources and sinks. For planar networks, they are able to reduce this exponential dependency to the minimum number of faces that cover all sources and sinks. For networks with a single source and a single edge with nonlinear cost function, Guisewite and Pardalos (1993) as well as Klinz and Tuy (1993) give polynomial-time algorithms. Tuy, Ghannadan, Migdalas, and Värbrand (1995) further isolate the dependency on the two critical parameters. They find an algorithm that runs in polynomial time if the number of sources as well as the number of arcs with nonlinear cost functions is bounded (see also Tuy 2000).

Besides the attempts to compute a global minimizer, there is also work on finding local minima. Gallo and Sodini (1979) and Guisewite and Pardalos (1991) devise local searches to tackle the problem. These are particularly interesting when it comes to finding Nash equilibria in network cost-sharing games. More on this is postponed to Section 5.3.1.

Network cost-sharing games. In their seminal paper, Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008) examine network design from a game theoretic perspective and introduce network cost-sharing games. The authors show that the price of anarchy in this model with constant total edge cost is the number of players, even in the case of undirected broadcast games. Subsequently, they focus on the price of stability. While their focus lies on constant total edge cost, they extend their results to more general settings including nonincreasing player cost.

The price of stability. The first upper bound on the price of stability is obtained by Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008) via the **potential function method**. It bounds the cost of a Nash equilibrium that results from executing the improving dynamics on a social optimum. The achieved upper bound of $H(n)$ is based on relating the social cost and potential function value of strategy profiles. The authors show that this bound is best possible for directed networks and constant total edge cost (even in the case of broadcast games). The following applies to undirected networks with constant total edge cost unless noted otherwise. Kawase and Makino (2012, 2013) examine the social cost of global potential minimizers, which are a subset of all Nash equilibria. The ratios of the maximal and minimal cost of such a potential minimizer to the cost of a social optimum are called **potential-optimal price of anarchy** and **potential-optimal price of stability**, respectively. The authors show

that both these prices are in $\mathcal{O}(\sqrt{\log n})$ and $\Omega(\sqrt{\log \log n})$ for broadcast games. Dissler, Feldmann, Klimm, and Mihalák (2013, 2015) take a similar approach for the general setting by analyzing the social cost of a global potential minimizer. They end up with the upper bound $(1 - \Theta(n^{-4}))H(n)$. Mamageishvili, Mihalák, and Montemezzani (2018, 2014) refine this method by analyzing strategy profiles obtained from combining a global potential minimizer with a social optimum. This gives an upper bound of $H(n/2) + \varepsilon$ for any $\varepsilon > 0$ and large enough n .

Better bounds on the price of stability are only known when the players share a common source. By exploiting the tree structure of a Nash equilibrium, Li (2009) refines the analysis of the potential function method and proves an upper bound of $\mathcal{O}(\log n / \log \log n)$ on the price of stability in multicast games. For broadcast games, a series of papers culminates in a constant price of stability. The evolved method is referred to as the **homogenization-absorption framework**. Its foundation is laid by Fiat, Kaplan, Levy, Olonetsky, and Shabo (2006). Just like the potential function method, their algorithm starts with a social optimum and transforms it into an equilibrium through improving moves. These moves, however, are chosen carefully such that the strategy profile stays relatively close to the initial social optimum. Players with similar strategies in the social optimum also have similar strategies throughout the improving steps, which is later dubbed **homogenization**. Due to this homogenization, when a player deviates, other players that are close in the social optimum have an incentive to follow this deviation. This effect is used in a process called **absorption**. These ideas yield an upper bound of $\mathcal{O}(\log \log n)$. Lee and Ligett (2013) introduce additional moves to the improving dynamics, in which groups of players deviate. Thereby, the cost of some deviating players might increase, but the potential function still decreases. This leads to a better homogenization and absorption. A charging scheme based on a preorder traversal of the social optimum results in an upper bound of $\mathcal{O}(\log \log \log n)$. Finally, Bilò, Flammini, and Moscardelli (2013, 2020) achieve a constant upper bound. They leverage that every vertex is a player to improve significantly on the homogenization. Their charging scheme uses the partition of edges into classes of exponentially growing cost. In an effort to transfer these results to more general settings, Freeman, Haney, and Panigrahi (2016) extend the homogenization-absorption framework to the class of multicast games on quasi-bipartite networks. These are multicast games in which every edge is incident to at least one player’s sink (or the common source). The authors obtain again a constant upper bound on the price of stability in this class.

In all cases, there is a large gap between the upper and lower bounds on the price of stability. In the general case, Christodoulou, Chung, Ligett, Pyrga, and van Stee (2009) and Bilò, Caragiannis, Fanelli, and Monaco (2013) give families of instances which yield lower bounds of approximately 1.826 and 2.245. The best lower bound for multicast games is 1.862 by Bilò, Caragiannis, Fanelli, and Monaco (2010, 2013). The same authors optimize the weights in an instance by Fiat, Kaplan, Levy, Olonetsky, and Shabo (2006) to raise their lower bound from approximately 1.714 to 1.818 for broadcast games.

As determining the price of stability has proven a tough problem, there is research on instances with a small number of players and networks with ring topology. For two players, Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008)

and Christodoulou, Chung, Ligett, Pyrga, and van Stee (2009) show a price of stability of $4/3$ in multicast and general networks, respectively. For three players, the work of Bilò and Bove (2011), Bilò, Caragiannis, Fanelli, and Monaco (2010, 2013), Christodoulou, Chung, Ligett, Pyrga, and van Stee (2009), and Disser, Feldmann, Klimm, and Mihalák (2015) culminates in the intervals $[1.571, 1.634]$ and $[1.524, 1.532]$ for the price of stability in general games and multicast games, respectively, and a value of around 1.485 for the price of stability in broadcast games. The fact that even for three players the exact value is not known illustrates the difficulty of the problem. In networks with ring topology, the price of stability is determined to be $3/2$ in general games by Fanelli, Leniowski, Monaco, and Sankowski (2012, 2015) and $4/3$ in multicast games by Mamageishvili, Mihalák, and Montemezzani (2018, 2014).

Computational complexity. A first result that suggests the hardness of computing Nash equilibria is by Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008). They construct an instance with constant total edge cost on which improving dynamics simulates a binary counter when initialized with a specific strategy profile. Hence, the running time of improving dynamics can be exponential in the number of players n .

Syrkkanis (2010) shows that finding a Nash equilibrium is PLS-hard for nonuniform network cost-sharing games with specific cost function classes. He uses a tight reduction from the maximum cut problem. This result covers the case of uniform directed network games with constant total edge cost. Bilò, Flammini, Monaco, and Moscardelli (2015, 2021) extend the reduction to also apply to undirected networks with constant total edge cost.

Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden (2004, 2008) show that computing a Nash equilibrium of minimal social cost in directed multicast games is NP-hard via a reduction from 3D matching. Syrgkanis (2010) shows the same hardness for undirected multicast games. Chekuri, Chuzhoy, Lewin-Eytan, Naor, and Orda (2006, 2007) prove that it is NP-hard to compute a global potential minimizer in undirected multicast games via a reduction from a variant of the satisfiability problem based on a result by Lund and Yannakakis (1993, 1994).

On the positive side, Albers and Lenzner (2010, 2013) show that any social optimum in an undirected multicast game is a $H(n)$ -approximate Nash equilibrium. For fixed $\alpha > 0$, Bilò, Flammini, Monaco, and Moscardelli (2015, 2021) give a polynomial-time algorithm to compute an $\Omega(n^{\alpha-1})$ -approximate equilibrium for network cost-sharing games with cost function $f(k) = k^\alpha$.

Several variants and extensions of network cost-sharing games are present in the literature. We only include a subset of them.

Sequential games. One line of research treats a variant in which players enter the game sequentially and choose a best response at the moment of entry to the present players. In this setting, Charikar, Karloff, Mathieu, Naor, and Saks (2008) prove the bounds

$\mathcal{O}(\log^2 n)$ and $\Omega(\log n)$ on the price of anarchy in multicast games. When a subsequent phase of best-response dynamics is added to obtain a Nash equilibrium, Chekuri, Chuzhoy, Lewin-Eytan, Naor, and Orda (2006, 2007) show that the price of anarchy in multicast games is bounded by $\mathcal{O}(\sqrt{n} \log^2 n)$ and $\Omega(\log n / \log \log n)$. Charikar, Karloff, Mathieu, Naor, and Saks (2008) improve these to $\mathcal{O}(\log^3 n)$ and $\Omega(\log n)$. Further, they examine the result of interleaving the player entries with best-response deviations. Mamageishvili, Mihalák, and Montemazzani (2018, 2014) treat sequential multicast games on ring topologies.

Relaxed and refined equilibria. Lee and Ligett (2013) examine a solution concept which they call go-it-alone equilibrium. It is a relaxation of the Nash equilibrium as it assumes that players only want to deviate if it is cheaper for them to pay their new strategy all on their own. Under this relaxation the price of stability in multicast games is constant for constant total edge cost.

Another line of research considers a refinement of Nash equilibria, the so-called strong Nash equilibria. Those are Nash equilibria in which no coalition of players can deviate and improve on the cost of every involved player. Epstein, Feldman, and Mansour (2007, 2009) show the existence of strong Nash equilibria in some network topologies. Then they show that the price of anarchy for this notion is bounded by $H(n)$ for general cost functions (in games that admit a strong Nash equilibrium). Albers (2008, 2009) refines this bound for approximate strong Nash equilibria and proves a lower bound of $\Omega(\sqrt{\log n})$.

Weighted players. In the model as examined in this chapter, all players using an edge share its cost equally. In a variation, every player is assigned a weight and the cost of an edge is shared proportionally to the weights of the present players. H. Chen and Roughgarden (2009, 2006) show that already in games with three weighted players Nash equilibria may fail to exist. Subsequently, they study approximate equilibria in this setting. Albers and Lenzner (2010) analyze to what extent social optima represent approximate Nash equilibria for weighted players. Chekuri, Chuzhoy, Lewin-Eytan, Naor, and Orda (2006) deal with the possible nonexistence of Nash equilibria in a different way. They show that equilibria always exist in a fractional version of network cost-sharing games with weighted players.

5.1.2 Our Contribution

We examine Nash equilibria of network cost-sharing games in terms of their computational complexity and their efficiency. Most existing work on these games focuses on the cost function $f(k) = 1/k$. We provide a wider angle, by considering general positive, nonincreasing functions f such that $k \mapsto kf(k)$ is nondecreasing and concave.

In terms of the complexity, our main results fortify the intractability of finding Nash equilibria. We discuss formulations for finding Nash equilibria as local optimization problems and relate them to the minimum concave-cost flow problem. These formulations yield natural algorithmic approaches, which, however, do not generally run in polynomial

time under common assumptions in complexity theory. This hardness is revealed by several results. We extend the work of Bilò, Flammini, Monaco, et al. (2015, 2021) and Syrgkanis (2010) to obtain PLS-hardness for all nonconstant cost function f . Further, we show that it is NP-hard to compute specific Nash equilibria even in broadcast games, specifically Nash equilibria of minimal social cost and global minimizers of the potential function. Finally, we give an example that provides evidence for slow convergence of the improving dynamics in multicast games.

In terms of the efficiency of Nash equilibria, we exactly determine the price of anarchy for arbitrary cost functions f as $f(1)/f(n)$ and improve the bounds on the price of stability. An improved analysis of the potential function method by Anshelevich et al. (2004, 2008) yields the better upper bound of $1 + \ln(f(1)/f(n))$ on the price of stability for general cost functions f . For the classes of affine total edge cost $kf(k) = 1 + s(k - 1)$ and polynomial total edge cost $kf(k) = k^\alpha$, we obtain the even tighter bounds $1 + W_0((1 - s)/(se))$ and $1/\alpha$, respectively. (Here, W_0 denotes the Lambert W function.) For broadcast games, we advance the homogenization-absorption framework by Bilò, Flammini, and Moscardelli (2013, 2020), Fiat et al. (2006), and Lee and Ligett (2013). We are able to significantly simplify the algorithm at its core and the charging scheme used for its analysis. This facilitates a constant upper bound on the price of stability of 265 for broadcast games with constant total edge cost, which is a magnitude smaller than (the estimate of) the best previously existing bound. Further, we obtain constant upper bounds for broadcast games in which $kf(k)$ is bounded. We augment these results by investigating the lower bounds on the price of stability. The structure of worst-case instances highly depends on the cost function f . Hence, we focus on fan graphs which provide the currently best lower bounds of 20/11 for broadcast games with constant total edge cost as found by Bilò, Caragiannis, et al. (2010, 2013) and Fiat et al. (2006). We prove that the known bound is the best possible within a natural class of instances. In addition, we obtain lower bounds for broadcast games with affine and polynomial total edge cost.

5.2 Structure of Nash equilibria

This section repeats and establishes basic results on the structure of Nash equilibria. The foundation for these results and large parts of the analysis in later sections is provided by the potential function. Rosenthal (1973) finds that it reflects the change of the total player cost in a unilateral deviation, as stated in the next theorem. We include its short proof due to its fundamental importance.

Theorem 5.1 (Potential game). *Let $\sigma \in \Sigma$ be a strategy profile and let $\hat{\sigma} \in \Sigma$ result from a player $i \in [n]$ unilaterally deviating in σ . Then it holds*

$$\Phi(\hat{\sigma}) - \Phi(\sigma) = C_i(\hat{\sigma}) - C_i(\sigma).$$

Proof. As $\hat{\sigma}$ differs from σ only by a unilateral deviation of a player $i \in [n]$, we know

$$n_{\hat{\sigma}}(e) - n_{\sigma}(e) = \begin{cases} 1 & \text{if } e \in \hat{\sigma}_i \setminus \sigma_i \\ -1 & \text{if } e \in \sigma_i \setminus \hat{\sigma}_i \\ 0 & \text{otherwise.} \end{cases}$$

Applying this to the definitions of the potential function and total player cost, we obtain

$$\begin{aligned} \Phi(\hat{\sigma}) - \Phi(\sigma) &= \sum_{e \in \hat{\sigma}_i \setminus \sigma_i} c_e f(n_{\hat{\sigma}}(e)) - \sum_{e \in \sigma_i \setminus \hat{\sigma}_i} c_e f(n_{\sigma}(e)) \\ &= \sum_{e \in \hat{\sigma}_i} c_e f(n_{\hat{\sigma}}(e)) - \sum_{e \in \sigma_i} c_e f(n_{\sigma}(e)) = C_i(\hat{\sigma}) - C_i(\sigma). \quad \blacksquare \end{aligned}$$

An immediate consequence of [Theorem 5.1](#) is that the set of local minimizers of the potential is exactly the set of Nash equilibria. This is discussed further in [Section 5.3.1](#).

The contribution of zero-cost edges to the total player cost (and the potential function) is zero. Hence, they behave slightly different than edges with positive cost. To deal with them, we make the following basic observation.

Observation 5.2 (Zero-cost edges). *Let $E' \subseteq E$ such that $c_e = 0$ for all $e \in E'$. Set $G' = G/E'$ to be the graph G after contracting the edges in E' . Then the total player costs of $\sigma \in \Sigma$ in the cost-sharing game on G are exactly the total player costs of σ/E' in the game on G' , that is, $C_i(\sigma) = C_i(\sigma/E')$ for all $i \in [n]$. In particular, their social costs are the same, and, σ is a Nash equilibrium in G if and only if σ/E' is a Nash equilibrium in G' .*

This observation essentially allows us to separate edges with zero and positive cost. We obtain our first structural result for Nash equilibria in network cost-sharing games.

Lemma 5.3 (Support of Nash equilibria). *Let $E_0 := \{e \in E \mid c_e = 0\}$ be the set of zero edges. For every Nash equilibrium $\sigma' \in \Sigma$ of a network cost-sharing game with n players and strictly decreasing cost function f , there is a Nash equilibrium $\sigma \in \Sigma$ such that*

- (i) $\sigma_i \setminus E_0 = \sigma'_i \setminus E_0$ for all $i \in [n]$, and
- (ii) the set of edges $\bigcup_{i \in [n]: v \in V(\sigma_i)} \sigma_i$ is a tree for every vertex $v \in V$.

Proof. We show that the support of σ' is a tree up to edges of cost zero. To break ties on E_0 , we need the function $\hat{\Phi}(\sigma) := \sum_{e \in E} H(n_{\sigma}(e))$. Note that $\hat{\Phi}$ is the potential function if we assume unit total edge cost. For all edges $e \in E$, define the corresponding cost $\hat{c}_{\sigma}(e) := 1/n_{\sigma}(e)$ if $n_{\sigma}(e) > 0$ and $\hat{c}_{\sigma}(e) = +\infty$ otherwise. Further, denote by $\hat{C}_i(\sigma)$ the total player cost of player $i \in [n]$ in $\sigma \in \Sigma$ with respect to \hat{c}_{σ} .

Let $\sigma \in \Sigma$ be such that it fulfills (i) and $\hat{\Phi}(\sigma)$ is minimal under all such choices. [Lemma 5.3](#) implies that σ is a Nash equilibrium. To see (ii), fix an arbitrary $v \in V$. Let $I := \{i \in [n] \mid v \in V(\sigma_i)\}$ be the set of players whose strategy is incident to v . Let T be a lexicographically shortest path tree from the source v with respect to c_{σ} and \hat{c}_{σ} ; that

is, for all $w \in V$ the path $T[v, w]$ minimizes $\widehat{c}_\sigma(T[v, w])$ under all shortest paths with respect to c_σ . We finish the proof by showing that the strategies of all players $i \in I$ lie in T .

Assume that there is $i \in I$ such that $\sigma_i \neq T[s_i, t_i]$. Set $\widehat{\sigma}$ to be the strategy profile that results from i unilaterally deviating in σ to $\widehat{\sigma}_i = T[s_i, t_i]$. Then,

$$c_\sigma(\widehat{\sigma}_i) \leq c_\sigma(T[s_i, v]) + c_\sigma(T[v, t_i]) \leq c_\sigma(\sigma_i) = C_i(\sigma)$$

holds, because σ_i is split by v into an s_i - v path and a v - t_i path.

If $\sigma_i \setminus E_0 \neq T[s_i, t_i] \setminus E_0$, then either $\widehat{\sigma}_i \setminus E_0 \subsetneq \sigma_i \setminus E_0$ or $(\widehat{\sigma}_i \setminus \sigma_i) \setminus E_0 \neq \emptyset$ has to be true. In the former case, clearly $C_i(\widehat{\sigma}) < C_i(\sigma)$ contradicts σ being a Nash equilibrium. Otherwise, f being strictly decreasing yields the contradiction

$$C_i(\widehat{\sigma}) = c_\sigma(\widehat{\sigma}_i \cap \sigma_i) + c'_\sigma(\widehat{\sigma}_i \setminus \sigma_i) < c_\sigma(\widehat{\sigma}_i) \leq C_i(\sigma).$$

Hence, $\sigma'_i \setminus E_0 = \sigma_i \setminus E_0 = \widehat{\sigma}_i \setminus E_0 = T[s_i, t_i] \setminus E_0$. Consequently, $\widehat{\Phi}(\sigma) \leq \widehat{\Phi}(\widehat{\sigma})$ follows from the choice of σ . Similarly to the preceding, we consider the two cases $\widehat{\sigma}_i \subsetneq \sigma_i$ and $\widehat{\sigma}_i \setminus \sigma_i \neq \emptyset$. In the first case, we get the contradiction $\widehat{\Phi}(\widehat{\sigma}) < \widehat{\Phi}(\sigma)$. In the second case, we repeat the preceding arguments and obtain

$$\widehat{C}_i(\widehat{\sigma}) = \widehat{c}_\sigma(\widehat{\sigma}_i \cap \sigma_i) + \widehat{c}'_\sigma(\widehat{\sigma}_i \setminus \sigma_i) < \widehat{c}_\sigma(\widehat{\sigma}_i) \leq \widehat{c}_\sigma(T[s_i, v]) + \widehat{c}_\sigma(T[v, t_i]) \leq \widehat{c}_\sigma(\sigma_i) \leq \widehat{C}_i(\sigma),$$

which is a contradiction based on [Theorem 5.1](#). Thus, in total $\sigma_i = T[s_i, t_i]$ holds for all $i \in I$ and σ fulfills (ii) as well. \blacksquare

For multicast and broadcast games, [Lemma 5.3](#) implies the following structural result as also found by Fiat et al. (2006).

Lemma 5.4 (Support of Nash equilibria in multicast games). *Let $E_0 = \{e \in E \mid c_e = 0\}$. For every Nash equilibrium $\sigma' \in \Sigma$ of a multicast game with strictly decreasing cost function f , there is a Nash equilibrium $\sigma \in \Sigma$ such that $\sigma_i \setminus E_0 = \sigma'_i \setminus E_0$ for all players i and $\text{supp}(\sigma)$ is a tree. In a broadcast game, $\text{supp}(\sigma)$ is a spanning tree.*

Proof. The statement is a corollary of [Lemma 5.3](#). The strategies of all players are incident to r . Thus, [Lemma 5.3](#) in particular implies that $\text{supp}(\sigma)$ is a tree. As in a broadcast game every vertex is a player connecting to r , it is spanning for those. \blacksquare

[Lemma 5.4](#) shows that Nash equilibria in multicast and broadcast games with strictly decreasing cost function are supported by a tree up to zero-cost edges. In particular, if all edges have positive cost, every Nash equilibrium is supported by a tree. For an equilibrium that is not supported by a tree there is an equilibrium on a tree of the very same social cost. [Lemma 5.4](#) enables us in many cases to restrict the attention to tree-supported equilibria. Therefore, we discuss them in terms of trees.

Consider a multicast game with root r . For a Nash equilibrium $\sigma \in \Sigma$ that is supported by a tree S , the strategy σ_v of a player $v \in V$ is the unique r - v path in S , which we denote by $S[r, v]$. Hence, S fully determines σ . We also say that S **induces** σ . This extends

to all trees that span the set of all players and the root. Also the difference between the strategies of two players $u, v \in V$ can be interpreted in terms of S . We have that $\sigma_v \setminus \sigma_u$ is exactly the path in S from v to the vertex where $S[v, r]$ and $S[u, r]$ meet. The latter is the lowest common ancestor of u and v in S and we get $\sigma_v \setminus \sigma_u = S[v, \text{lca}(u, v)]$.

Their tree structure leads to a particularly nice characterization of Nash equilibria in multicast and broadcast games. It is sufficient to consider deviations which are the concatenation of a path in $E \setminus S$ and a path in S .

Theorem 5.5 (Nash equilibria in multicast games). *Let $\sigma \in \Sigma$ be a strategy profile in a multicast game that is supported by a tree $S = \text{supp}(\sigma)$. Then σ is a Nash equilibrium if and only if it holds for all vertices $u, v \in V(S)$ and all u - v paths $P \subseteq E \setminus S$ that are internally vertex disjoint to $V(S)$ that*

$$c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) \leq c(P)f(1) \quad (\text{NE})$$

where the notation $\sigma_w = S[r, w]$ is extended to all $w \in V$.

Proof. Assume σ is a Nash equilibrium. Fix $u, v \in V(S)$ and a u - v path $P \subseteq E \setminus S$ which does not intersect S . Let $\hat{\sigma}$ be the strategy profile that results from v unilaterally deviating in σ to $\hat{\sigma}_v = \sigma_u \cup P$. Applying the definition of a Nash equilibrium to σ yields

$$\begin{aligned} 0 &\leq C_v(\hat{\sigma}) - C_v(\sigma) \\ &= c_\sigma((\sigma_u \cup P) \cap \sigma_v) + c'_\sigma((\sigma_u \cup P) \setminus \sigma_v) - c_\sigma(\sigma_v) \\ &= c'_\sigma(\sigma_u \setminus \sigma_v) + c(P)f(1) - c_\sigma(\sigma_v \setminus \sigma_u). \end{aligned}$$

For the converse direction, assume that σ is not a Nash equilibrium. Then there is a player $w \in V$ who can decrease her cost by deviating. Choose $\hat{\sigma}_w$ to be an improving move of w in σ that minimizes $|\hat{\sigma}_w \setminus S|$. Without loss of generality, we can assume that $\hat{\sigma}_w$ is a path. As σ is supported by a tree, it follows that the improving move $\hat{\sigma}_w$ contains at least one edge that is not in S . Define $P = \hat{\sigma}_w[u, v]$ by $u, v \in V(S)$ such that $S[r, u] = \hat{\sigma}_w[r, u]$, $V(P) \cap V(S) = \{u, v\}$, and $P \neq \emptyset$; that is, P is the inclusion-maximal subpath of $\hat{\sigma}_w$ that is closest to r and internally vertex disjoint to $V(S)$. We compare to $\hat{\sigma}_w$ the strategy $\sigma_v \cup \hat{\sigma}_w[v, w]$, which uses σ_v instead of $\sigma_u \cup P$. Observe that $\sigma_v \cup \hat{\sigma}_w[v, w]$ contains strictly less edges in $E \setminus S$. Hence, it cannot be improving for w in σ as this would contradict the choice of $\hat{\sigma}_w$. Thus, we get

$$\begin{aligned} 0 &< C_w(\sigma_{-w}, \sigma_v \cup \hat{\sigma}_w[v, w]) - C_w(\sigma_{-w}, \hat{\sigma}_w) \\ &\leq c_\sigma(\sigma_v \setminus (\sigma_u \cup P)) - c'_\sigma((\sigma_u \cup P) \setminus \sigma_v) \\ &= c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) - c(P)f(1). \quad \blacksquare \end{aligned}$$

Corollary 5.6 (Nash equilibria in broadcast games). *Let $\sigma \in \Sigma$ be a strategy profile in a broadcast game that is supported by a tree $S = \text{supp}(\sigma)$. Then σ is a Nash equilibrium if and only if it holds for all edges $e = \{u, v\} \in E \setminus S$ that*

$$c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) \leq c_e f(1).$$




Figure 5.2: Braess's paradox for the price of stability in an instance with $f(k) = 1/k$ (left). The social optima and Nash equilibria are highlighted in yellow and violet, respectively. The deletion of $\{r, 3\}$ decreases the price of stability from $(22 + 2\varepsilon)/(17 + \varepsilon)$ to 1 (right).

Ultimately, the sparsity of the support of Nash equilibria is based on the concavity of the potential function. As we assume that $k \mapsto kf(k)$ is concave as well, we obtain that social optima are supported by trees as well (up to edges of cost zero). This also results from the following observation on pairs of related games.

Observation 5.7. *Let $f: \mathbb{N} \rightarrow \mathbb{R}$ be a cost function. Then the potential function of the game with cost function f is equal to the social cost function of the game with cost function $\bar{f}: \mathbb{N} \rightarrow \mathbb{R}, k \mapsto F(k)/k$. Hence, a social optimum with respect to \bar{f} is a global potential minimizer with respect to f and, hence, a Nash equilibrium.*

We finish this section, by exploring the effect of changing the network on the set of Nash equilibria and social optima. Clearly, a social optimum is preserved when deleting edges which are not in its support. We observe that the same is true for Nash equilibria.

Observation 5.8 (Nash equilibria under edge deletion). *Let $\sigma \in \Sigma$ be a Nash equilibrium in a network cost-sharing game and let $E' \subseteq E$ be a subset of edges that contains $\text{supp}(\sigma)$. Then σ is also a Nash equilibrium in the network cost-sharing game restricted to E' .*

Crucially, [Observation 5.8](#) deals only with deleting edges that are not part of the equilibrium. Clearly, deleting arbitrary edges can destroy Nash equilibria. It turns out, that deleting edges can also create Nash equilibria. We obtain an effect similar to the paradox by Braess ([1968](#)). In the example illustrated in [Figure 5.2](#), the price of stability decreases on deleting an edge. The cost function of the example is $f(k) = 1/k$. The edges are annotated with their costs c_e . We compare the price of stability before and after deleting the edge $\{r, 3\}$. First of all, a social optimum is given by the path $r-1-2-3$ with a social cost of $17 + \varepsilon$. The induced strategy profile by this path becomes a Nash equilibrium after deleting $\{r, 3\}$. Before, however, the best Nash equilibrium is induced by the tree $\{\{r, 2\}, \{1, 2\}, \{2, 3\}\}$. Its social cost evaluates to $22 + 2\varepsilon$.

Algorithm 3: Improving dynamics

Input: $G = (V, E), r \in V, c \in \mathbb{R}_{\geq 0}^E$, player cost function f **Output:** Nash equilibrium $\sigma \in \bar{\Sigma}$

```

1 Choose an arbitrary strategy profile  $\sigma \in \Sigma$ ;
2 while  $\exists i \in [n], \sigma'_i \in \Sigma_i: C_i(\sigma_{-i}, \sigma'_i) < C_i(\sigma)$  do
3   |  $\sigma \leftarrow (\sigma_{-i}, \sigma'_i)$ ;
4 end
5 return  $\sigma$ ;

```

5.3 Computational Complexity of Nash Equilibria

In this section, we examine the complexity of computing Nash equilibria. On the positive side, we discuss formulations which lead to algorithmic approaches. On the negative side, we give hardness results.

5.3.1 Formulations

Given the definition of a Nash equilibrium, it is not clear how to construct one algorithmically. The potential function allows to apply methods from (local) optimization. We compare the standard approach of the improving dynamics to a local search that stems from a concave-cost flow formulation. Further, we discuss a relaxation of this flow formulation.

Improving dynamics. [Theorem 5.1](#) suggests a simple algorithm for computing a Nash equilibrium. An arbitrary initial strategy profile is gradually manipulated. As long as there is a player who wants to deviate from the current strategy profile, let her unilaterally defect to an improving strategy. This algorithm is called **improving dynamics**. A formal description is presented in [Algorithm 3](#). Its finite convergence is an immediate consequence of [Theorem 5.1](#). Indeed, the potential decreases strictly by each unilateral deviation. As there is only a finite number of strategy profiles, this process has to terminate after a finite number of steps.

In multicast games, [Lemma 5.4](#) shows that Nash equilibria are essentially supported by trees. We obtain a variant of the improving dynamics which keeps the tree structure of a strategy profile throughout the process.

Lemma 5.9 (Improving tree moves in multicast games). *Let σ be a strategy profile in a multicast game that is supported by a tree $S = \text{supp}(\sigma) \subseteq E$. If there are vertices $u, v \in V(S)$ and a u - v path P that is internally vertex disjoint to $V(S)$ and violates (NE) then the strategy profile $\hat{\sigma}$ induced by the tree $S \setminus \{e_S(v)\} \cup P$ fulfills $\Phi(\hat{\sigma}) < \Phi(\sigma)$ and $C_w(\hat{\sigma}) \leq C_w(\sigma_{-w}, \hat{\sigma}_w) < C_w(\sigma)$ for every $w \in D_S(v)$.*

Proof. Let $u, v \in V(S)$ and the u - v path P that is internally vertex disjoint to $V(S)$ violate (NE), that is, $c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) > c(P)f(1)$. Note that $u \notin D_S(v)$ as

Algorithm 4: Improving tree dynamics

Input: $G = (V, E), r \in V, c \in \mathbb{R}_{\geq 0}^E$, player cost function f

Output: Nash equilibrium $\sigma \in \Sigma$

- 1 Let σ be a strategy profile induced by an arbitrary spanning tree;
 - 2 **while** $\exists u, v \in V(\text{supp}(\sigma)), u$ - v path P internally vertex disjoint to $\text{supp}(\sigma)$:
 $c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) > c(P)f(1)$ **do**
 - 3 | Update σ to be the induced strategy profile of $S \setminus \{e_S(v)\} \cup P$;
 - 4 **end**
 - 5 **return** σ ;
-

otherwise $c_\sigma(\sigma_v \setminus \sigma_u) = 0$ contradicts the strict inequality. Hence, $S \setminus \{e_S(v)\} \cup P$ is indeed a tree. Let $\hat{\sigma}$ be its induced strategy profile. As all $w \in D_S(v)$ share the path σ_v and $\hat{\sigma}_v$ in σ and $\hat{\sigma}$, respectively, we have $\hat{\sigma}_w \setminus \sigma_w = \hat{\sigma}_v \setminus \sigma_v$ and $\sigma_w \setminus \hat{\sigma}_w = \sigma_v \setminus \hat{\sigma}_v$. Therefore, we get

$$\begin{aligned} C_w(\sigma_{-w}, \hat{\sigma}_w) - C_w(\sigma) &= c'_\sigma(\hat{\sigma}_w \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \hat{\sigma}_w) \\ &= c'_\sigma(\sigma_u \setminus \sigma_v) + c(P)f(1) - c_\sigma(\sigma_v \setminus \sigma_u) < 0. \end{aligned}$$

Due to the tree structure of σ and $\hat{\sigma}$, we know for every $e \in E$ and $w \in D_S(v)$

$$n_{\hat{\sigma}}(e) - n_\sigma(e) = \begin{cases} |D_S(v)| & \text{if } e \in \hat{\sigma}_v \setminus \sigma_v, \\ -|D_S(v)| & \text{if } e \in \sigma_v \setminus \hat{\sigma}_v, \\ 0 & \text{otherwise} \end{cases} = |D_S(v)|(n_{\sigma_{-w}, \hat{\sigma}_w}(e) - n_\sigma(e)).$$

Note that $v \in D_S(v)$ and so $|D_S(v)| \geq 1$. Thus, we get for the player cost of all $w \in D_S(v)$

$$\begin{aligned} C_w(\hat{\sigma}) - C_w(\sigma) &= \sum_{e \in \hat{\sigma}_w \setminus \sigma_w} c_e f(n_{\hat{\sigma}}(e)) - \sum_{e \in \sigma_w \setminus \hat{\sigma}_w} c_e f(n_\sigma(e)) \\ &\leq \sum_{e \in \hat{\sigma}_w \setminus \sigma_w} c_e f(n_\sigma(e) + 1) - \sum_{e \in \sigma_w \setminus \hat{\sigma}_w} c_e f(n_\sigma(e)) = C_w(\sigma_{-w}, \hat{\sigma}_w) - C_w(\sigma). \end{aligned}$$

Finally, the change of the potential is

$$\begin{aligned} \Phi(\hat{\sigma}) - \Phi(\sigma) &= \sum_{e \in E} c_e (F(n_{\hat{\sigma}}(e)) - F(n_\sigma(e))) \\ &\leq |D_S(v)| \sum_{e \in \hat{\sigma}_w \setminus \sigma_w} c_e f(n_\sigma(e) + 1) - |D_S(v)| \sum_{e \in \sigma_w \setminus \hat{\sigma}_w} c_e f(n_\sigma(e)) \\ &= |D_S(v)|(C_w(\sigma_{-w}, \hat{\sigma}_w) - C_w(\sigma)) \\ &< 0. \end{aligned} \quad \blacksquare$$

Based on [Lemma 5.9](#), we devise a variant of the improving dynamics on the set of strategy profiles which are supported by trees. We call this the **improving tree dynamics**. Its pseudo code is listed in [Algorithm 4](#). If a profile that is supported by a tree is not in equilibrium, there is a whole set of players given by a subtree who want to deviate. Note that every single of these players benefits in this move. These players would even deviate successively one by one. Therefore, a move in the improving tree dynamics corresponds to a sequence of moves in the improving dynamics.

Local search. Improving dynamics can be viewed as local search. For a strategy profile $\sigma \in \Sigma$, we define its **unilateral neighborhood** as the set of strategy profiles that differ from σ by the strategy of exactly one player, that is,

$$N^{\text{unilateral}}(\sigma) = \{\sigma' \in \Sigma \mid \exists! i \in [n] \text{ such that } \sigma'_i \neq \sigma_i\}.$$

Then we obtain from [Theorem 5.1](#) the following characterization of Nash equilibria.

Corollary 5.10 (Nash equilibria as local minima). *The Nash equilibria in a network cost-sharing game are exactly the local minima of Φ with respect to the unilateral neighborhood $N^{\text{unilateral}}$.*

Hence, finding a Nash equilibrium is equivalent to computing a local minimum of the potential function. In particular, the global minimizer of the potential function is a Nash equilibrium. The latter can be formulated in terms of a flow problem. To do so, we need an orientation of the undirected graph $G = (V, E)$. Let $\vec{G} = (V, A)$ be the graph on the same vertex set as G such that every edge is replaced by two opposite arcs. We extend the cost to \vec{G} in the natural way, that is, $c_{v,w} = c_{\{v,w\}}$ for all $(v, w) \in A$.

$$\begin{aligned} \min \quad & \sum_{\substack{e \in E \\ e = \{v,w\}}} c_e F \left(\sum_{i \in [n]} x_{v,w}^{(i)} + \sum_{i \in [n]} x_{w,v}^{(i)} \right) & (\text{MCCFP}) \\ \text{s. t.} \quad & \sum_{a \in \delta^-(v)} x_a^{(i)} - \sum_{a \in \delta^+(v)} x_a^{(i)} = \begin{cases} 1 & \text{if } v = t_i \\ -1 & \text{if } v = s_i \\ 0 & \text{otherwise} \end{cases} & \text{for all } v \in V, i \in [n] \\ & x_a^{(i)} \in \mathbb{N} & \text{for all } a \in A, i \in [n] \end{aligned}$$

Strictly speaking, (MCCFP) is not a flow problem. Its objective function is not separable with respect to the total flow on every arc. Instead it involves the sum of the total flows on opposite arcs. This issue can be remedied by replacing two opposite arcs with a gadget as depicted in [Figure 5.3](#). The central arc of the gadget gets the cost of the original arc while the other arcs get cost zero. In that spirit, we can treat (MCCFP) as an (uncapacitated) multicommodity minimum concave-cost flow problem. In case of a multicast game with players $W \subseteq V$ this multicommodity flow problem simplifies to a




Figure 5.3: Two opposite arcs (left) are replaced by a gadget (right). The arcs are annotated with their cost.

singlecommodity flow problem.

$$\begin{aligned}
 & \min \sum_{a \in A} c_a F(x_a) && \text{(SSMCCFP)} \\
 & \text{s. t. } \sum_{a \in \delta^-(v)} x_a - \sum_{a \in \delta^+(v)} x_a = \begin{cases} 1 & \text{if } v \in W \\ 0 & \text{otherwise} \end{cases} && \text{for all } v \in V \setminus \{r\} \\
 & && x_a \in \mathbb{N} && \text{for all } a \in A
 \end{aligned}$$

Note that for the single commodity case, the objective does not need to consider the sum of the total flows on opposite arcs. This is due to the fact that in any minimum only one of two opposite arcs can carry flow. This stems from a directed version of [Lemma 5.4](#). Thus, (SSMCCFP) gives a formulation as (uncapacitated) single-source minimum concave-cost flow problem.

The extreme points of the uncapacitated single-source flow polyhedron correspond exactly to flows that are supported by an arborescence. These in turn correspond exactly to strategy profiles which are supported by trees. Gallo and Sodini (1979) show that two flows that are supported by arborescences are adjacent in (the one skeleton of) the polyhedron if and only if the union of their supports contains exactly one undirected cycle. Augmenting along this cycle generates one flow from the other. Subsequently, Gallo and Sodini (1979) propose a local search for solving uncapacitated single-source minimum concave-cost flow problems with respect to the polyhedral neighborhood

$$N^{\text{polyhedral}}(\sigma) = \{ \sigma' \in \Sigma \mid \text{supp}(\sigma) \cup \text{supp}(\sigma') \text{ contains a unique cycle} \}.$$

This local search is closely related to the improving tree dynamics. However, there is a subtle difference between them. The improving tree move guaranteed by [Lemma 5.9](#) is indeed adjacent to the initial strategy profile with respect to $N^{\text{polyhedral}}$. Hence, a local minimum with respect to $N^{\text{polyhedral}}$ is a Nash equilibrium. But not every potential decreasing move with respect to $N^{\text{polyhedral}}$ must give rise to a deviation in which every deviating player benefits. This is caused by the fact that local search requires the decrease of the potential only, whereas the improving dynamics requires the decrease of all deviating players' individual costs.

Nonlinear optimization. Finally, we consider another concept of local minima. In continuous nonlinear optimization a point is a local minimum if it assumes the smallest objective function value within a small (continuous) neighborhood. We relax (MCCFP) by dropping the integrality constraints on the flow. Let \widehat{F} be a continuous extension of F onto $\mathbb{R}_{\geq 0}$, that is, F and \widehat{F} agree on \mathbb{N} . Using this extension of the objective function to the relaxed feasible region, we obtain a nonlinear continuous optimization problem (MCCFP'). x is a local minimum of (MCCFP') if there is $\varepsilon > 0$ such that $F(x) \leq F(x')$ holds for every feasible point x' of (MCCFP') with $\|x' - x\|_{\infty} \leq \varepsilon$. Note that the specific choice of the uniform norm is not important, as the definition is invariant under different choices. The next results establish properties of local minima of (MCCFP') in dependence on the interpolation \widehat{F} .

Lemma 5.11. *Let \widehat{F} be a strictly concave, continuous extension of F and $c_a > 0$ for all $a \in A$. Then every local minimum of (MCCFP') is integer.*

Proof. In (MCCFP'), a strictly concave objective function is minimized over a polyhedron. Hence, any local optimum is attained at a vertex of the polyhedron. As the flow polyhedron is integer, the statement follows. \blacksquare

Lemma 5.12. *Let \widehat{F} be a nondecreasing, concave, continuous extension of F . Then every Nash equilibrium corresponds to a local minimum of (MCCFP').*

Proof. Define $\widehat{\Phi}: \mathbb{R}^A \rightarrow \mathbb{R}, x \mapsto \sum_{\{v,w\} \in E} c_e \widehat{F}(x_{v,w} + x_{w,v})$. Let σ be a Nash equilibrium, and let x be the flow induced by σ . We prove that x has minimal objective value under all feasible points within distance $\frac{1}{n}$ with respect to the uniform norm. To this end, let x' be a feasible point of (MCCFP') such that $\|x' - x\|_{\infty} \leq \frac{1}{n}$. Denote by \mathcal{P}_i the set of s_i - t_i paths in G , and by \mathcal{C} the set of cycles. Then $nx' - (n-1)x$ is also feasible to (MCCFP'). Decomposing this flow into elementary flows on paths and cycles yields $y_P^{(i)} \geq 0$ for $i \in [n], P \in \mathcal{P}_i$ with $\sum_{P \in \mathcal{P}_i} y_P^{(i)} = 1$ for all $i \in [n]$ and $y_C \geq 0$ for $C \in \mathcal{C}$ such that

$$nx' - (n-1)x = \sum_{i \in [n]} \sum_{P \in \mathcal{P}_i} y_P^{(i)} \mathbb{1}_P + \sum_{C \in \mathcal{C}} y_C \mathbb{1}_C.$$

For a player $i \in [n]$ and a path $P \in \mathcal{P}_i$, the flow $x - \mathbb{1}_{\sigma_i} + \mathbb{1}_P$ corresponds to the strategy profile that results from i deviating in σ to P . As σ is a Nash equilibrium, we know

$$\widehat{\Phi}(x) \leq \sum_{i \in [n]} \sum_{P \in \mathcal{P}_i} \frac{y_P^{(i)}}{n} \widehat{\Phi}(x - \mathbb{1}_{\sigma_i} + \mathbb{1}_P).$$

We apply concavity of $\widehat{\Phi}$ to the right-hand side and obtain

$$\begin{aligned}\widehat{\Phi}(x) &\leq \widehat{\Phi}\left(\sum_{i \in [n]} \sum_{P \in \mathcal{P}_i} \frac{y_P^{(i)}}{n} (x - \mathbb{1}_{\sigma_i} + \mathbb{1}_P)\right) \\ &= \widehat{\Phi}\left(x - \frac{1}{n} \sum_{i \in [n]} \mathbb{1}_{\sigma_i} + x' - \frac{n-1}{n}x - \sum_{C \in \mathcal{C}} \frac{y_C}{n} \mathbb{1}_C\right) \\ &= \widehat{\Phi}\left(x' - \sum_{C \in \mathcal{C}} \frac{y_C}{n} \mathbb{1}_C\right).\end{aligned}$$

Finally, $\widehat{\Phi}(x) \leq \widehat{\Phi}(x')$ follows, because \widehat{F} is nondecreasing. \blacksquare

Lemma 5.13. *Let \widehat{F} be the linear interpolation of F . Then every local minimum of (MCCFP') is a Nash equilibrium.*

Proof. Let $x \in \mathbb{R}^A$ be a local minimum of (MCCFP'); that is, there exists $\varepsilon > 0$ such that $\widehat{\Phi}(x) \leq \widehat{\Phi}(x')$ for all feasible $x' \in \mathbb{R}^E$ with $\|x' - x\|_\infty \leq \varepsilon$. By Lemma 5.11, x is integer and, hence, corresponds to a strategy profile $\sigma \in \Sigma$. For a player $i \in [n]$ and a path $P \in \mathcal{P}_i$, let $\widehat{\sigma}$ be the resulting strategy profile when i deviates in σ to the strategy P , and let $\widehat{x} \in \mathbb{R}^A$ be the corresponding flow. Then $\|x - \widehat{x}\|_\infty \leq 1$ shows that \widehat{F} is linear on the segment between x and \widehat{x} . Therefore, we get

$$\widehat{\Phi}(x) \leq \widehat{\Phi}((1 - \varepsilon)x + \varepsilon\widehat{x}) = \widehat{\Phi}(x) + \varepsilon(\widehat{\Phi}(\widehat{x}) - \widehat{\Phi}(x)).$$

Consequently, $\Phi(\sigma) = \widehat{\Phi}(x) \leq \widehat{\Phi}(\widehat{x}) = \Phi(\widehat{\sigma})$. Because $\widehat{\sigma}$ emerged from an arbitrary unilateral deviation, this shows that σ is a Nash equilibrium. \blacksquare

Unfortunately, the linear interpolation of F is inherently nonsmooth. Lemma 5.13 still holds for continuous extensions of F that are close to the linear interpolation. Yet, such an extension does not seem to yield a smooth function which would be beneficial for nonlinear optimization.

5.3.2 Intractability Results

In this section, we give lower bounds on the complexity of computing Nash equilibria in network cost-sharing games. For arbitrary structure of the sources and sinks, we show that finding a Nash equilibrium for any fixed strictly decreasing edge cost function is PLS-complete. The considered neighborhood is that of unilateral deviations. In other words, the improving dynamics does not run in polynomial time unless $P = \text{PLS}$. Note that the same is true for the improving tree dynamics as any sequence of improving tree moves translates to a sequence of improving moves which is longer by a polynomial factor. The reduction for PLS-hardness does not work for multicast games. This is why we proceed by examining the complexity of computing specific Nash equilibria. We

prove that it is NP-hard to compute a Nash equilibrium which minimizes the social cost in broadcast games for any fixed strictly decreasing cost function. Further, computing the global minimizer of the potential function, which is also a Nash equilibrium, is NP-hard for broadcast games with fixed strictly decreasing cost function.

5.3.2.1 PLS-Hardness of Computing Nash Equilibria

Syrkkanis (2010) shows PLS-hardness of computing Nash equilibria in nonuniform network cost-sharing games for a certain class of cost functions. Bilò, Flammini, Monaco, et al. (2015, 2021) extend this hardness to uniform games with cost function $f(k) = 1/k$. Both results are based on a reduction from the maximum cut problem via a specific subclass of cost-sharing games. We refer to this class as **intermediate cost-sharing games**. Our reduction uses a similar construction and shows PLS-hardness for uniform network cost-sharing games with arbitrary fixed nonconstant cost function f .

Theorem 5.14 (PLS-completeness). *Let $f: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be a cost function such that $f(2) < f(1)$. Then computing a Nash equilibrium in a network cost-sharing game with cost function f is PLS-complete (with respect to the neighborhood of unilateral deviation).*

For the proof of this theorem, we use the PLS-hardness of a subclass of (general) cost-sharing games as shown by Syrgkanis (2010). A cost-sharing game is defined by its players, resources, the players' strategies, and the resources' cost functions. The sets of players and resources can both be arbitrary (finite) sets. In this context, a strategy is a subset of resources. The cost function of a resource maps the number of players using it to the cost each player has to pay for using it. The cost of a player is then the sum of the cost of all resources she uses. For network cost-sharing games the resources are given by the edges in the network.

An instance of the intermediate cost-sharing game has a particular structure. There are $n \in \mathbb{N}$ players. For every (unordered) pair of players $\{i, j\}$ there are two resources $r_{i,j}^0$ and $r_{i,j}^1$. Hence, the set of resources is given by $\mathcal{R} := \{r_{i,j}^k \mid 1 \leq i < j \leq n, k \in \{0, 1\}\}$. Every player $i \in [n]$ has exactly two strategies which are

$$\begin{aligned} \sigma_i^0 &:= \{r_{0,i}^0, \dots, r_{i-1,i}^0, r_{i,i+1}^1, \dots, r_{i,n}^1\} \text{ and} \\ \sigma_i^1 &:= \{r_{0,i}^1, \dots, r_{i-1,i}^1, r_{i,i+1}^0, \dots, r_{i,n}^0\}. \end{aligned}$$

Note that a resource $r_{i,j}^k$ is only part of a strategy of the players i and j . Further, all strategies of i and j contain exactly one of $r_{i,j}^0$ and $r_{i,j}^1$. In particular, every resource can be used by at most two players. Hence, we need to specify the resource cost functions only for one and two players. For every $1 \leq i < j \leq n$, the resource cost function of $r_{i,j}^0$ and $r_{i,j}^1$ is defined by a weight $w_{i,j} \in \mathbb{R}_{\geq 0}$. If a player exclusively uses the resource $r_{i,j}^0$ or $r_{i,j}^1$, she pays $w_{i,j}$ for it. Resources used by more than one player do not incur any cost.

Proof of Theorem 5.14. We transform an instance of the intermediate cost-sharing game to an instance of a network cost-sharing game. Then we show the correspondence of their respective sets of equilibria.

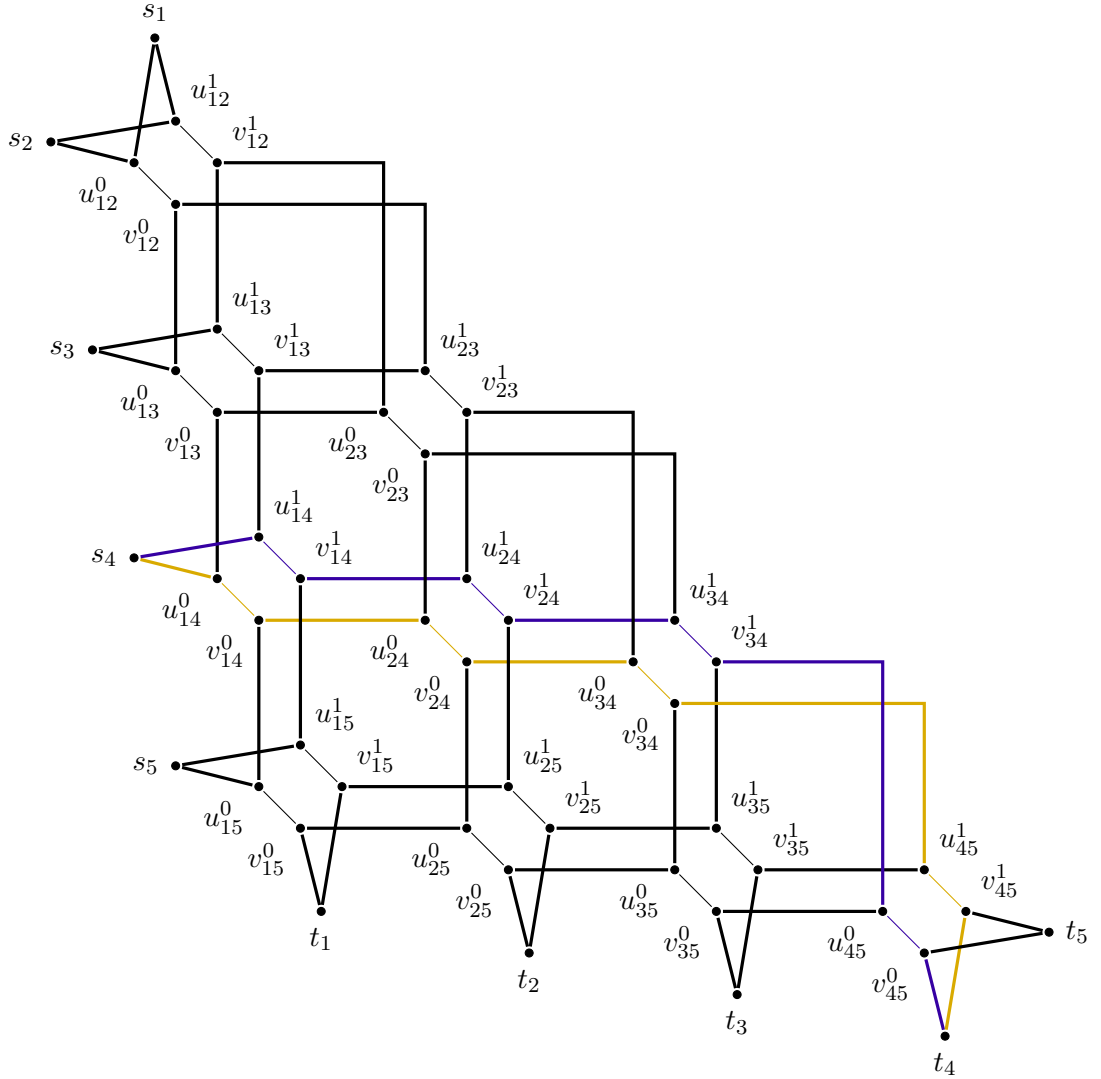


Figure 5.4: An instance resulting from an intermediate cost-sharing game with $n = 5$ players. Thick connections represent paths of length L . Thin edges correspond to the resources of the intermediate game. The two highlighted paths are the dominant strategies of player 4, which correspond to her strategies in the intermediate game.

The network cost-sharing game. Let $n \in \mathbb{N}$ be the number of players in the intermediate cost-sharing game. Further, let $r_{i,j}^0$ and $r_{i,j}^1$ be the resources with weight $w_{i,j} \in \mathbb{R}_{\geq 0}$ of the pair of players $1 \leq i < j \leq n$. The constructed network is illustrated in Figure 5.4. Every resource $r_{i,j}^k$ is represented by an edge $\{u_{i,j}^k, v_{i,j}^k\}$ with cost function $k \mapsto w_{i,j} f(k) / (f(1) - f(2))$. This function is chosen such that the difference between the player cost for one and two players is exactly $w_{i,j}$.

Each thick connection in Figure 5.4 represents an undirected path of length $L \in \mathbb{N}$. All edges of such a path are assigned the cost function Wf , where $W \in \mathbb{R}_{\geq 0}$ is chosen an order

of magnitude larger compared to the remaining costs in the instance. Therefore, these paths are called **heavy paths**. Heavy paths are used to restrict the strategies which can appear in an equilibrium. To achieve that the heavy paths' cost is essentially constant and does not depend too much on the number of players using them, we introduce dummy players. For every of the L edges of a heavy path, $D \in \mathbb{N}$ dummy players are added and associated with the edge. Their source and sink are the endpoints of their associated edge. The length L is chosen carefully to ensure that for every dummy player the associated edge is a dominant strategy.

For every player $i \in [n]$ of the intermediate cost-sharing game, there is a player with source s_i and sink t_i in the network cost-sharing game. Due to the high cost of the heavy paths, every such player naturally minimizes the number of such paths in her strategy. Only as a secondary objective, the players minimize the cost incurred by edges that correspond to resources of the intermediate cost-sharing game. As a result, in an equilibrium player i plays one of the two paths

$$P_i^k := s_i - u_{1,i}^k - v_{1,i}^k - \dots - u_{i-1,i}^k - v_{i-1,i}^k - u_{i,i+1}^{1-k} - v_{i,i+1}^{1-k} - \dots - u_{i,n}^{1-k} - v_{i,n}^{1-k} - t_i$$

for $k \in \{0, 1\}$. Note that these dominant strategies pairwise intersect only in edges which correspond to resources of the intermediate cost-sharing game. For $i \in [n], k \in \{0, 1\}$, the strategy P_i^k corresponds to the strategy σ_i^k of the intermediate cost-sharing game. The edges in P_i^k that correspond to resources of the intermediate game are exactly σ_i^k .

The parameters. We choose the parameters L, D , and W and show subsequently that they yield the desired properties. First, we show that if n is large enough, there is $D \in [2n^3]$ satisfying $nf(D+1) < (n + \frac{1}{2})f(D+n)$. Assume for a contradiction that this is not the case. Then, in particular, it holds

$$f(l(n-1)+2) \leq \frac{2m}{2n+1} f((l-1)(n-1)+2) \quad \text{for all } l \in [n(2n+1)].$$

We chain all these inequalities and use that $k \mapsto kf(k)$ is nondecreasing to obtain

$$\frac{1}{n(2n+1)(n-1)+2} \leq \frac{f(n(2n+1)(n-1)+2)}{f(1)} \leq \left(1 - \frac{1}{2n+1}\right)^{(2n+1)n} \leq e^{-n}.$$

This is a contradiction for large enough n , as the exponential function grows faster than any polynomial. Therefore, a D with the desired property exists. Set $L := 2(D+1)$ and $W := n(n-1) \max_{1 \leq i < j \leq n} w_{i,j} / (f(1) - f(2))$.

The strategies of dummy players in equilibria. First, we prove that in a Nash equilibrium σ every dummy player plays her associated edge. Assume there is a heavy path P such that there are associated dummy players of P which do not play their associated edge. By [Lemma 5.3](#), we may assume without loss of generality that all dummy players associated with the same edge have the same strategy in σ . Let $E' \subseteq P$ be the set of edges in P that are not played by their associated dummy players. Note that all such dummy players use all edges of P except for their associated one. Further,

let K be the number of players using P which are not associated with it. Every such external player either uses all edges of P or none. We obtain that

$$n_\sigma(e) = \begin{cases} K + (|E'| - 1)D & \text{for } e \in E' \\ K + (|E'| + 1)D & \text{for } e \in P \setminus E'. \end{cases}$$

Choose d to be one of the dummy players associated with an edge in E' . We compare σ to the strategy profile $\hat{\sigma}$ which results from σ by d unilaterally deviating to play her associated edge. If $|E'| \geq 2$, then $L \geq 3$ and f being nonincreasing yield

$$\begin{aligned} C_d(\hat{\sigma}) &= Wf\left(K + (|E'| - 1)D + 1\right) \\ &< (L - |E'|)Wf\left(K + (|E'| + 1)D\right) + (|E'| - 1)Wf\left(K + (|E'| - 1)D\right) \\ &\leq C_d(\sigma). \end{aligned}$$

As this contradicts σ being a Nash equilibrium, $|E'|$ must be equal to one. But then $k \mapsto kf(k)$ being nondecreasing implies

$$\begin{aligned} C_d(\hat{\sigma}) &= Wf(K + 1) \leq W\frac{K + 2D}{K + 1}f(K + 2D) \leq 2DWf(K + 2D) \\ &< (L - 1)Wf(K + 2D) \leq C_d(\sigma) \end{aligned}$$

which again contradicts the choice of σ . It follows that all dummy players use their associated edge in a Nash equilibrium.

The strategies of other players in equilibria. Next, we show that a player $i \in [n]$ plays either P_i^0 or P_i^1 if all dummy players play their associated edge. Assume the latter is the case for a strategy profile σ . From the preceding, we know that every edge is used by at least D and at most $D + n$ players. Note that P_i^0 or P_i^1 contain exactly n heavy paths, and any other s_i - t_i path contains at least $n + 1$ heavy paths. Assume there is $i \in [n]$ such that $\sigma_i \notin \{P_i^0, P_i^1\}$. Let $\hat{\sigma}$ be the strategy profile that agrees with σ except for $\hat{\sigma}_i$ being P_i^0 . Then our choice of the parameters gives

$$\begin{aligned} C_i(\hat{\sigma}) &\leq (n - 1) \max_{j \in [n]} \frac{w_{ij}}{f(1) - f(2)} f(1) + nLWf(D + 1) \\ &< \frac{W}{n} f(1) + \left(n + \frac{1}{2}\right) LWf(D + n) \\ &\leq \frac{W}{n} f(1) - W\frac{2(D + 1)}{2(D + n)} f(1) + (n + 1)LWf(D + n) \\ &\leq (n + 1)LWf(D + n) \leq C_i(\sigma). \end{aligned}$$

This shows that player i wants to deviate from σ .

Correspondence of equilibria. Let σ be a strategy profile such that all dummy players play their associated edge and all players $i \in [n]$ play one of their paths P_i^0 or P_i^1 . We know that all Nash equilibria in the network cost-sharing game are of this form. Let the strategy profile $\hat{\sigma}$ in the intermediate cost-sharing game be defined as follows. Player $i \in [n]$ plays $\hat{\sigma}_i = \sigma_i^k$ if and only if i plays $\sigma_i = P_i^k$ for $k \in \{0, 1\}$. As the paths P_i^k are pairwise disjoint with respect to the heavy paths and contain exactly the edges that correspond to the resources in σ_i^k , we obtain

$$\begin{aligned} C_i(\sigma) &= nLWf(D+1) + \sum_{\substack{j \in [n] \setminus \{i\}: \\ \sigma_i \cap \sigma_j = \emptyset}} w_{i,j} \frac{f(1)}{f(1) - f(2)} + \sum_{\substack{j \in [n] \setminus \{i\}: \\ \sigma_i \cap \sigma_j \neq \emptyset}} w_{i,j} \frac{f(2)}{f(1) - f(2)} \\ &= nLWf(D+1) + \sum_{j \in [n] \setminus \{i\}} w_{i,j} \frac{f(2)}{f(1) - f(2)} + \sum_{\substack{j \in [n] \setminus \{i\}: \\ \sigma_i \cap \sigma_j = \emptyset}} w_{i,j} \\ &= nLWf(D+1) + \sum_{j \in [n] \setminus \{i\}} w_{i,j} \frac{f(2)}{f(1) - f(2)} + C_i(\hat{\sigma}). \end{aligned}$$

where we set $w_{i,j} = w_{j,i}$ for all $i \neq j$. We see that $C_i(\sigma)$ is increasing in $C_i(\hat{\sigma})$. If σ is an equilibrium, so must be $\hat{\sigma}$. This finishes the proof. \blacksquare

Note that the PLS-reduction in the proof of [Theorem 5.14](#) is tight. The set of relevant strategy profiles in the constructed network cost-sharing game is given by all profiles in which dummy players play their associated edge and all players $i \in [n]$ play either P_i^0 or P_i^1 .

The construction does not seem to generalize to multicast games or even broadcast games. The possibility to pick individual source and sink vertices for all players is used in two ways. On the one hand, the dummy players are bound to their associated edge. On the other hand, [Lemma 5.4](#) shows that a common source would result in equilibria which are supported by trees. This appears insufficient to model the rich interaction patterns between players in the intermediate cost-sharing game.

Remark 5.15. [Theorem 5.14](#) can be extended to all cost functions which are not constant. If there is $K \in \mathbb{N}$ such that $f(K+2) < f(K+1)$, all edges $\{u_{ij}^k, v_{ij}^k\}$ for $1 \leq i < j \leq n$ and $k \in \{0, 1\}$ can be replaced by a path of length $2(K+1)$. Further, every edge of these paths is associated with additional K dummy players. As for the dummy players on heavy paths of the construction, also for these new dummy players their dominant strategy is their associated edge. This results in the nondummy players either meeting K or $K+1$ other players instead of zero or one other player on these edges. Thus, the difference $f(K+1) - f(K+2)$ appear instead of $f(1) - f(2)$ in the analysis.

Note that an equilibrium can be efficiently computed in the case $f(k) = 1$. Then there is no benefit for the players to share edges, hence a strategy profile is in equilibrium if and only if every player chooses a shortest path with respect to c .




Figure 5.5: The broadcast game resulting from an exact 3-set cover problem with universe $U = \{u_1, \dots, u_6\}$ and family of subsets $\mathcal{S} = \{S_1, \dots, S_4\}$. The strategy profile supported by the highlighted edges corresponds to the exact covering $\{S_2, S_4\}$.

5.3.2.2 NP-Hardness of Computing a Minimum-Cost Nash Equilibrium

We show that it is NP-hard to compute a Nash equilibrium which minimizes the social cost in broadcast games with fixed strictly decreasing edge cost function f . This was previously only known for multicast games with constant total edge cost. We devise a reduction from the exact 3-set cover problem. Let an instance be given by the universe U and the family of subsets \mathcal{S} . Define the following broadcast game. The graph $G = (V, E)$ is given by the set of vertices $V := \{r\} \cup \mathcal{S} \cup U$ where r is a distinct root vertex. The edge set is given by $E := \{\{r, S\} \mid S \in \mathcal{S}\} \cup \{\{S, u\} \mid u \in S \in \mathcal{S}\} \cup \{\{r, u\} \mid u \in U\}$ with weights

$$c_e = \begin{cases} 1 + \frac{f(4)}{f(1)} & \text{if } \exists u \in U: e = \{r, u\}, \text{ and} \\ 1 & \text{otherwise.} \end{cases}$$

See [Figure 5.5](#) for an illustration.

Theorem 5.16 (Correspondence of packings and equilibria). *The set of packings of (U, \mathcal{S}) corresponds one-to-one to the set of equilibria in G . Moreover, a packing $\mathcal{P} \subseteq \mathcal{S}$ corresponds to an equilibrium with social cost $f(1)|\mathcal{S} \setminus \mathcal{P}| + f(4)|\mathcal{P}| + (f(1) + f(4))|U|$.*

Proof. We define a canonical injection from packings to equilibria. A subsequent examination of the equilibria's structure implies its surjectivity.

Mapping packings to equilibria. Let $\mathcal{P} \subseteq \mathcal{S}$ be a packing of U . Define the strategy profile $\sigma^{\mathcal{P}}$ by

$$\sigma_v^{\mathcal{P}} := \begin{cases} \{\{r, S\}, \{S, v\}\} & \text{if } \exists S \in \mathcal{P}: v \in S, \text{ and} \\ \{\{r, v\}\} & \text{otherwise.} \end{cases}$$

The player cost of $v \in V$ evaluates to

$$C_v(\sigma^{\mathcal{P}}) = \begin{cases} f(1) & \text{if } v \in \mathcal{S} \setminus \mathcal{P}, \\ f(4) & \text{if } v \in \mathcal{P}, \text{ and} \\ f(1) + f(4) & \text{if } v \in U. \end{cases}$$

It can be checked that $\sigma^{\mathcal{P}}$ is an equilibrium. Its social cost is as claimed.

The height of equilibria. Let σ be a Nash equilibrium. Due to [Lemma 5.4](#), we may assume that $T := \text{supp}(\sigma)$ is a tree rooted at r . We claim that the height of T is two. Assume that it is larger. Let $u \in V$ be a leaf of maximal depth in T . If $u \in \mathcal{S}$, then $|\sigma_u| > 1$ and $C_u(\sigma) > f(1)$. In that case, deviating to $\{u, r\}$ would result in a cost of at most $f(1)$ for u and, hence, be an improving move. As σ is a Nash equilibrium, however, it follows that $u \in U$. Let $S \in \mathcal{S}$ be the parent of u in T . Since we assumed $|\sigma_u| > 2$, S has at most two children in T and $C_u(\sigma) > f(1) + f(3)$. Hence, u wants to deviate to $\{u, r\}$ which incurs a cost of at most $f(1) + f(4)$. This contradicts σ being a Nash equilibrium. It follows that T has height 2.

Surjectivity. Let $u \in U$ such that $|\sigma_u| = 2$ and let $S \in \mathcal{S}$ be the parent of u in T . As $|S| = 3$, S has at most three children in T . Hence, it follows $n_\sigma(\{S, r\}) \leq 4$. For the player cost of u , we obtain

$$C_u(\sigma) = f(n_\sigma(\{u, S\})) + f(n_\sigma(\{S, r\})) \geq f(1) + f(4).$$

As f is strictly decreasing, the inequality holds with equality if and only if $n_\sigma(\{S, r\}) = 4$. This means that u wants to deviate to $\{r, u\}$ in the case that $n_\sigma(\{S, r\}) < 4$. In total, we get that $n_\sigma(\{S', r\}) \in \{1, 4\}$ for all $S' \in \mathcal{S}$. Consequently, we get the correspondence $\sigma = \sigma^{\mathcal{P}}$ for the packing $\mathcal{P} := \{S' \in \mathcal{S} \mid n_\sigma(\{S', r\}) = 4\}$. ■

Corollary 5.17 (NP-hardness of minimum-cost Nash equilibria). *Let f be a strictly decreasing cost function. It is NP-hard to compute a Nash equilibrium of minimal social cost in undirected broadcast games with edge cost function f .*

Proof. An instance of the exact 3-set cover problem has an exact covering if and only if it has a packing of size $\frac{1}{3}|U|$. By [Theorem 5.16](#), this is the case if and only if the constructed broadcast game has a Nash equilibrium of social cost at most $f(1)|\mathcal{S}| + (\frac{2}{3}f(1) + \frac{4}{3}f(4))|U|$. ■

5.3.2.3 NP-Hardness of Computing a Global Potential Minimizer

The global minimizer of the potential is in particular a Nash equilibrium. This gives a natural approach to computing one. It is NP-hard, however, to compute this specific equilibrium even in broadcast games. The hardness is very much the same as the hardness of computing social optima for network cost-sharing games with arbitrary edge cost functions. This connection is based on [Observation 5.7](#). For completeness, we include the proof here.




Figure 5.6: Illustration of the reduction for the hardness of global potential minimization. The depicted broadcast game is based on an exact 3-set cover instance with universe $U = \{u_1, \dots, u_6\}$ and the family of subsets $\mathcal{S} = \{S_1, \dots, S_4\}$. The highlighted edges support a strategy profile that corresponds to the exact set covering $\{S_2, S_4\}$.

The reduction resembles that of the previous section. Again, we reduce from the exact 3-set cover problem. Fix a resource cost function f . Let an instance of exact 3-set cover be given by a universe U and a family of subsets $\mathcal{S} \subseteq 2^U$. We transform it into an instance of a broadcast game. The set of vertices is given by a root vertex r , all elements of the universe U , and all sets in \mathcal{S} , that is, $V = \{r\} \cup \mathcal{S} \cup U$. We add all edges between r and \mathcal{S} as well as all edges between \mathcal{S} and U . The cost of edges incident to r is set to 1, the cost of the other edges is set to $f(1)/f(3)$. The resulting instance is illustrated in Figure 5.6.

Theorem 5.18 (NP-hardness of the global potential minimizer). *Let f be strictly decreasing. It is NP-hard to compute a global minimizer of the potential function.*

Proof. Let $\sigma \in \Sigma$ be a global minimizer of the potential Φ in a network cost-sharing game resulting from the preceding reduction. Due to Theorem 5.1, we know that σ is a Nash equilibrium. By Lemma 5.4, the support of σ is a spanning tree. Assume that its height is larger than two. Let $u \in V$ be a leaf of maximal depth in this tree. If u is in \mathcal{S} its path to r has to start with an edge of cost $f(1)/f(3)$. It follows that $C_u(\sigma) > f(1)^2/f(3)f(1)$. But deviating to $\{r, u\}$ would cost $f(1)$ only, which contradicts σ being a Nash equilibrium. Hence, $u \in U$. Let $S \in \mathcal{S}$ be the parent of u in $\text{supp}(\sigma)$. Assume the strategy σ_u is not $\{\{r, S\}, \{S, u\}\}$. Then, S has at most two children in $\text{supp}(\sigma)$ and $C_u(\sigma) > (f(1) + f(3))f(1)/f(3)$. Therefore, u would deviate to $\{\{r, S\}, \{S, u\}\}$ which costs at most $(f(1)/f(3) + 1)f(1)$. In total, we get that $\text{supp}(\sigma)$ has height two. In particular, $\sum_{S \in \mathcal{S}} n_\sigma(\{r, S\}) = |\mathcal{S}| + |U|$.

Using the structure of σ , we obtain

$$\Phi(\sigma) = \sum_{S \in \mathcal{S}} F(n_\sigma(\{r, S\})) + \frac{f(1)}{f(3)} \sum_{u \in S \in \mathcal{S}} F(n_\sigma(\{S, u\})) = \sum_{S \in \mathcal{S}} F(n_\sigma(\{r, S\})) + \frac{f(1)^2}{f(3)} |U|.$$

Further, we know that $1 \leq n_\sigma(\{r, S\}) \leq 4$ holds for all $S \in \mathcal{S}$. Thus, we can bound the right-hand side with the help of the concavity of F and find

$$\begin{aligned}
 \Phi(\sigma) &\geq \sum_{S \in \mathcal{S}} \left(\frac{4 - n_\sigma(\{r, S\})}{3} F(1) + \frac{n_\sigma(\{r, S\}) - 1}{3} F(4) \right) + \frac{f(1)^2}{f(3)} |U| \\
 &= \sum_{S \in \mathcal{S}} \left(F(1) + \frac{F(4) - F(1)}{3} (n_\sigma(\{r, S\}) - 1) \right) + \frac{f(1)^2}{f(3)} |U| \\
 &= |\mathcal{S}| F(1) + \left(\sum_{S \in \mathcal{S}} n_\sigma(\{r, S\}) - |\mathcal{S}| \right) \frac{F(4) - F(1)}{3} + \frac{f(1)^2}{f(3)} |U| \\
 &= |\mathcal{S}| F(1) + |U| \left(\frac{F(4) - F(1)}{3} + \frac{f(1)^2}{f(3)} \right)
 \end{aligned}$$

As we assume that f is strictly decreasing, F is strictly concave. Hence, the inequality is tight if and only if $n_\sigma(\{r, S\}) \in \{1, 4\}$ for all $S \in \mathcal{S}$. The latter is equivalent to the existence of an exact covering of U . If $n_\sigma(\{r, S\}) \in \{1, 4\}$ for all $S \in \mathcal{S}$, then $\{S \in \mathcal{S} \mid n_\sigma(\{r, S\}) = 4\}$ is such an exact cover. An exact covering $\mathcal{C} \subseteq \mathcal{S}$ on the other hand induces a strategy profile $\sigma^{\mathcal{C}}$ with $\sigma_S^{\mathcal{C}} = \{r, S\}$ for all $S \in \mathcal{S}$ and $\sigma_u^{\mathcal{C}} = \{\{r, S\}, \{S, u\}\}$ for all $u \in S \in \mathcal{C}$. This profile clearly fulfills $n_{\sigma^{\mathcal{C}}}(\{r, S\}) \in \{1, 4\}$. It follows that computing a global potential minimizer is enough to decide on exact 3-set cover. \blacksquare

Note that an α -approximation to the global potential minimizer for some $\alpha \geq 1$ is not necessarily an α -approximate Nash equilibrium. This might seem to be the case on a first glance. While the change in the cost of a deviating player and the corresponding change of the potential are the same, the two notions of approximate solutions relate this value to different baselines. The absolute value of the potential function can be arbitrarily larger than the cost of a player.

5.3.2.4 Slowly Improving Dynamics

Improving dynamics provides a canonical algorithm for computing Nash equilibria. PLS-hardness in the general setting, however, shows that it can take exponential time to converge. The complexity of computing equilibria in multicast games is still open. We provide evidence that improving dynamics might converge slowly for multicast games with $f(k) = 1/k$ as well. Remember that the measure of progress for the dynamics is the potential function. If the potential function would decrease sufficiently in every step, the number of steps was bounded. We provide an example with unit constant total edge cost with a step of small improvement.

This is remarkable as Schäffer and Yannakakis (1991) observe that the exponential running time of local search algorithms stems from the weights and their binary encoding. It is not immediate, however, that network cost-sharing games with $c \equiv 1$ can be solved

in polynomial time. The reason is that in their case the weight of an edge e with k players is given by $c_e f(k)$ and not c_e alone.

Let $\sigma \in \Sigma$ be a strategy profile and $\widehat{\sigma}_i$ be an improving move of player $i \in [n]$. The change of the potential under this deviation is

$$\Phi(\sigma_{-i}, \widehat{\sigma}_i) - \Phi(\sigma) = \sum_{e \in \widehat{\sigma}_i \setminus \sigma_i} \frac{1}{n_\sigma(e) + 1} - \sum_{e \in \sigma_i \setminus \widehat{\sigma}_i} \frac{1}{n_\sigma(e)}.$$

To show that this change can be small, we give numbers $a_1, \dots, a_m \in \{-n, \dots, n\} \setminus \{0\}$ such that $\delta := \sum_{i=1}^m \frac{1}{a_i} < 0$ is exponentially small in n and m . The positive and negative a_i correspond to $n_\sigma(e) + 1$ for $e \in \widehat{\sigma}_i \setminus \sigma_i$ and $-n_\sigma(e)$ for $e \in \sigma_i \setminus \widehat{\sigma}_i$, respectively. Then δ is exactly the difference $\Phi(\sigma_{-i}, \widehat{\sigma}_i) - \Phi(\sigma)$.

Example 5.19 (Small potential decrease). Fix $1 \leq k < n$ and set $a_j = (-1)^{k+j-1}(n-j)$ for $j = 0, \dots, k$. We do the preceding calculation for the family of numbers where each a_j appears exactly $\binom{k}{j}$ times. In total, we get $\sum_{j=0}^k \binom{k}{j} = 2^k$ many numbers. We claim that

$$\delta := \sum_{j=0}^k \binom{k}{j} \frac{(-1)^{k+j-1}}{n-j} = -\frac{1}{(n-k)\binom{n}{k}}$$

Choosing k and n such that $n = 2^k$, then yields

$$\binom{n}{k} \geq \left(\frac{n}{\log n}\right)^{\log n} = n^{\log n - \log \log n} \geq n^{\frac{1}{2} \log n}$$

for large enough n . This shows that $\binom{n}{k}$ grows faster than any polynomial in n . Consequently, the absolute value of δ decreases faster than the inverse of any polynomial.

We finish by proving the claim with an induction on k . It can be checked quickly that it holds for $k = 1$. We assume that it holds for some $k \in \mathbb{N}$ and prove that it holds for $k + 1$ as well. By Pascal's rule and a subsequent index shift, we obtain

$$\sum_{j=0}^{k+1} \binom{k+1}{j} \frac{(-1)^{k+j}}{n-j} = \sum_{j=0}^k \binom{k}{j} \frac{(-1)^{k+j+1}}{n-1-j} + \sum_{j=0}^k \binom{k}{j} \frac{(-1)^{k+j}}{n-j}.$$

Applying the induction hypothesis to both sums yields

$$\sum_{j=0}^{k+1} \binom{k+1}{j} \frac{(-1)^{k+j}}{n-j} = -\frac{1}{(n-1-k)\binom{n-1}{k}} + \frac{1}{(n-k)\binom{n}{k}} = -\frac{1}{(n-k-1)\binom{n}{k+1}}.$$

This finishes the proof of the claim.

This example gives the algebraic basis for an improving step in a multicast game with very small potential decrease. The multicast game with a suitable strategy profile σ is depicted in [Figure 5.7](#). When player i deviates from the lower path σ_i to the upper path $\widehat{\sigma}_i$ the change of potential is exactly $\Phi(\sigma_{-i}, \widehat{\sigma}_i) - \Phi(\sigma) = -\frac{1}{(n-k)\binom{n}{k}}$. The example

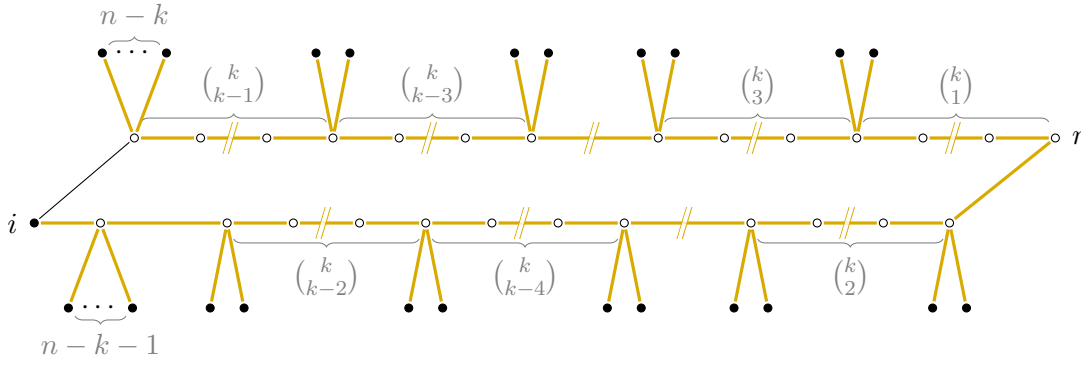


Figure 5.7: The deviation of player i from the highlighted strategy profile in this multicast game leads to a small decrease of the potential. Vertices with associated players are represented by filled nodes (\bullet). All total edge costs are constant and one.

shows that in the improving dynamics there can occur steps which decrease the potential only slightly. This is a necessary condition for the improving dynamics to take super-polynomially many steps to converge. It is an open question whether there are instances in which this actually happens.

Note that the numeric example cannot be transformed to an improving move in a broadcast game easily. The reason is that the same numbers appear several times. In a strategy profile of a broadcast game that is induced by a tree, however, the number of users on the edges increases strictly towards the root. In particular, every such number can appear at most once.

5.4 Efficiency of Nash Equilibria

In this section, we examine the increase of the social cost in network design that is caused by the selfish behavior of the players. We determine the price of anarchy exactly for network cost-sharing games with arbitrary fixed cost function f . Further, we obtain bounds on the price of stability. Firstly, we apply the potential function method in order to prove upper bounds for general structure of the sources and sinks. Secondly, we simplify the homogenization-absorption framework and improve its analysis to get a constant upper bound on the price of stability in broadcast games in which $kf(k)$ is bounded. Finally, we analyze a class of instances to obtain lower bounds for constant, affine, and polynomial total edge costs.

5.4.1 The Price of Anarchy

Anshelevich et al. (2008) show that the price of anarchy for network cost-sharing games with cost function $f(k) = 1/k$ is exactly n . We generalize the analysis to arbitrary fixed cost functions. The proof of the upper bound as well as the example for the lower bound are very similar to the known case. We include it for completeness.




Figure 5.8: A broadcast game with root r which exhibits a large price of anarchy. The edges are annotated with their cost c_e . The social optimum (yellow) and a Nash equilibrium of large social cost (violet) are highlighted.

Theorem 5.20 (The price of anarchy). *The price of anarchy for network cost-sharing games with n players and cost functions f is exactly $f(1)/f(n)$.*

Proof. We start by proving the upper bound on the price of anarchy, and continue by showing that it is tight by providing a class of instances.

Upper bound. Let $\sigma^* \in \Sigma$ be a social optimum and $\sigma \in \Sigma$ be a Nash equilibrium. As σ is an equilibrium, a player cannot improve by deviating from σ to her respective strategy in σ^* . We obtain $C_i(\sigma) \leq C_i(\sigma_{-i}, \sigma_i^*)$ for all $i \in [n]$. Using further that $f(n) \leq f(1)$, we get the inequality

$$C_i(\sigma) \leq C_i(\sigma_{-i}, \sigma_i^*) = \sum_{e \in \sigma_i^*} c_e f(n_{\sigma_{-i}, \sigma_i^*}(e)) \leq \sum_{e \in \sigma_i^*} c_e \frac{f(n_{\sigma^*}(e))}{f(n)} f(1) = \frac{f(1)}{f(n)} C_i(\sigma^*).$$

Summing up this inequality for every player $i \in [n]$ yields $f(n)C(\sigma) \leq f(1)C(\sigma^*)$.

Lower bound. To see that this bound is tight, consider the two highlighted strategy profiles for the instance depicted in Figure 5.8. It is a broadcast game with n vertices that want to connect to a root r . To do so, they can choose one of two parallel edges.

In the unique social optimum σ^* all players use the cheaper edge of total cost $f(n)$ (besides edges of cost zero). Its social cost is $C(\sigma^*) = nf(n)^2$. The strategy profile σ in which all players use the more expensive edge of total cost $f(1)$ (besides edges of cost zero) is a Nash equilibrium. Indeed, all player costs are $f(1)f(n)$, which is exactly the cost they would incur when unilaterally deviating to the cheaper edge. Hence, this instance has a price of anarchy of

$$\frac{C(\sigma)}{C(\sigma^*)} = \frac{nf(1)f(n)}{nf(n)^2} = \frac{f(1)}{f(n)}. \quad \blacksquare$$

Note that the proof of the upper bound on the price of anarchy fits the smoothness framework of Roughgarden (2009, 2015). In particular, network cost-sharing games are $(f(1)/f(n), 0)$ -smooth. Multicast and broadcast games have the same price of anarchy as the lower bound instance is in both these classes. We observe that the price of anarchy is constant for $n \rightarrow \infty$ if and only if f is bounded away from zero. The following lemma can be used to characterize this class of cost functions. An example are affine total edge costs $c_e k f(k) = c_e + s c_e (k - 1)$. Here, we obtain a price of anarchy of $1/s$.

Lemma 5.21. *Let $f: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be a cost function and $s \geq 0$. Then $f(k) \geq s$ holds for all $k \in \mathbb{N}$ if and only if $(k + 1)f(k + 1) - kf(k) \geq s$ for all $k \in \mathbb{N}$.*

Proof. We show $\lim_{k \rightarrow +\infty} f(k) = \lim_{k \rightarrow +\infty} (k + 1)f(k + 1) - kf(k)$. The statement then follows as f and $k \mapsto (k + 1)f(k + 1) - kf(k)$ are nonincreasing by the assumptions on f . Because both these functions are additionally nonnegative, we know that the limits exist.

Let $\hat{s} := \lim_{k \rightarrow +\infty} (k + 1)f(k + 1) - kf(k)$. Fix $\varepsilon > 0$. There is a $K \in \mathbb{N}$ such that

$$\hat{s} \leq (k + 1)f(k + 1) - kf(k) \leq \hat{s} + \varepsilon$$

holds for all $k \geq K$. It follows inductively that $(k - K)\hat{s} \leq kf(k) - Kf(K) \leq (k - K)(\hat{s} + \varepsilon)$ for all $k \geq K$. Dividing by k and taking the limit yields

$$\hat{s} = \lim_{k \rightarrow +\infty} \frac{(k - K)\hat{s}}{k} \leq \lim_{k \rightarrow +\infty} \frac{kf(k) - Kf(K)}{k} \leq \lim_{k \rightarrow +\infty} \frac{(k - K)(\hat{s} + \varepsilon)}{k} = \hat{s} + \varepsilon.$$

As ε was chosen arbitrarily, it follows $\lim_{k \rightarrow +\infty} f(k) = \lim_{k \rightarrow +\infty} f(k) - \frac{K}{k}f(K) = \hat{s}$. This finishes the proof. \blacksquare

5.4.2 Upper Bounds on the Price of Stability

We use two techniques for providing upper bounds to the price of stability. For network cost-sharing games with general source-sink structure, the potential function method by Anshelevich et al. (2004, 2008) is still asymptotically the best known upper bound. We improve its analysis for arbitrary fixed cost functions and even further for specific classes of cost functions. For broadcast games, we continue the work of Bilò, Flammini, and Moscardelli (2013, 2020), Fiat et al. (2006), and Lee and Ligett (2013) and simplify the homogenization-absorption framework significantly. This leads to an improved analysis and a constant which is by a magnitude smaller than the previously best bound.

5.4.2.1 The Potential Function Method

Anshelevich et al. (2008) devise a method for the proof of existence of Nash equilibria with low social cost. The key idea lies in relating the potential function value $\Phi(\sigma)$ and the social cost $C(\sigma)$ for a strategy profile $\sigma \in \Sigma$. Let $\sigma^* \in \Sigma$ be a social optimum. Initiating the improving dynamics (Algorithm 3) with σ^* leads to a Nash equilibrium $\sigma \in \Sigma$ with a potential that is at most the potential of σ^* . On the one hand, $C(\sigma) \leq \Phi(\sigma)$

follows directly from the definitions of Φ and C . On the other hand, we can bound any summand of $\Phi(\sigma^*)$ relative to the respective summand of $C(\sigma^*)$ and get

$$\begin{aligned} C(\sigma) \leq \Phi(\sigma) &\leq \Phi(\sigma^*) = \sum_{e \in E} c_e F(n_{\sigma^*}(e)) \\ &\leq \max_{\bar{n} \in [n]} \frac{F(\bar{n})}{\bar{n}f(\bar{n})} \sum_{e \in E} c_e n_{\sigma^*}(e) f(n_{\sigma^*}(e)) = \max_{\bar{n} \in [n]} \frac{F(\bar{n})}{\bar{n}f(\bar{n})} C(\sigma^*). \end{aligned}$$

Bounding the maximum on the right-hand side immediately gives an upper bound on the price of stability. Due to its reliance on the potential, this approach is called the **potential function method**. With it, Anshelevich et al. (2004, 2008) show that the n -th harmonic number $H(n)$ is an upper bound. We improve this bound for general uniform cost functions.

Theorem 5.22 (The price of stability). *The price of stability for a network cost-sharing game with n players and player cost function f is bounded by $1 + \ln(f(1)/f(n))$.*

Proof. We combine two bounds on f to estimate $F(\bar{n})/(\bar{n}f(\bar{n}))$ from above for $\bar{n} \in [n]$. First, we know $f(k) \leq f(1)$ for all $k \in \mathbb{N}$ because f is nonincreasing. Secondly, we assumed that $k \mapsto kf(k)$ is nondecreasing. Therefore, $f(k) \leq \bar{n}/k \cdot f(\bar{n})$ holds for all $1 \leq k \leq \bar{n}$. Both together yield

$$\frac{F(\bar{n})}{\bar{n}f(\bar{n})} = \frac{\sum_{k=1}^{\bar{n}} f(k)}{\bar{n}f(\bar{n})} \leq \sum_{k=1}^{\bar{n}} \min\left\{\frac{f(1)}{\bar{n}f(\bar{n})}, \frac{1}{k}\right\}.$$

Clearly, the minimum on the right-hand side is nonincreasing in k . This allows us to bound the sum by an integral and we obtain

$$\frac{F(\bar{n})}{\bar{n}f(\bar{n})} \leq \int_{k=0}^{\bar{n}} \min\left\{\frac{f(1)}{\bar{n}f(\bar{n})}, \frac{1}{k}\right\} dk = \int_{k=0}^{\frac{\bar{n}f(\bar{n})}{f(1)}} \frac{f(1)}{\bar{n}f(\bar{n})} dk + \int_{k=\frac{\bar{n}f(\bar{n})}{f(1)}}^{\bar{n}} \frac{1}{k} dk = 1 + \ln\left(\frac{f(1)}{f(\bar{n})}\right).$$

Note that the right-hand side is nondecreasing in \bar{n} . Hence, the maximum in the bound of the potential function method is attained at $\bar{n} = n$. This gives the desired bound on the price of stability. \blacksquare

Note that $\ln(k) \leq H(k) \leq \ln(k) + 1$ holds for all $k \in \mathbb{N}$. Therefore, the bound of Theorem 5.22 essentially interpolates between the previously known bounds of 1 for linear total edge cost and $H(n)$ for constant total edge cost. Under the presence of multiple sinks and multiple sources, no better asymptotic upper bound for the price of stability is known.

The analysis of the previous theorem can be improved under the assumption that the total cost of edges grows at a certain rate.

Theorem 5.23 (The price of stability under increasing total cost). *Let $h: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be a nonincreasing function. The price of stability of a network cost-sharing game with cost function f such that $f(1) = h(1)$ and $(k+1)f(k+1) - kf(k) \geq h(k+1)$ for all $k \in \mathbb{N}$ is at most $1 + W_0((f(1) - h(n))/(eh(n)))$. Here, W_0 denotes (the principal branch of) the Lambert W function.*

Proof. We fix $\bar{n} \in [n]$. Since $k \mapsto kf(k)$ is concave, we know $(k+1)f(k+1) - kf(k) \geq h(k)$ for all $1 \leq k < \bar{n}$. This can be used to relate $f(k)$ to $f(\bar{n})$ by the telescoping sum

$$\bar{n}f(\bar{n}) - kf(k) = \sum_{l=k+1}^{\bar{n}} (lf(l) - (l-1)f(l-1)) \geq (\bar{n} - k)h(\bar{n}).$$

On the one hand, it follows $f(\bar{n}) \geq h(\bar{n})$ by setting $k = 1$ and using $f(1) = h(1) \geq h(\bar{n})$. On the other hand, rearranging this inequality leads to an upper bound on $f(k)$. As f is nonincreasing, we additionally can use the upper bound $f(k) \leq f(1)$ and get

$$F(\bar{n}) = \sum_{k=1}^{\bar{n}} f(k) \leq \sum_{k=1}^{\bar{n}} \min \left\{ f(1), h(\bar{n}) + \frac{\bar{n}}{k} (f(\bar{n}) - h(\bar{n})) \right\}.$$

Again, we use the fact that the minimum on the right-hand side is nonincreasing in k . This allows us to bound the sum by an integral, which yields

$$\begin{aligned} F(\bar{n}) &\leq \int_{k=0}^{\bar{n}} \min \left\{ f(1), h(\bar{n}) + \frac{\bar{n}}{k} (f(\bar{n}) - h(\bar{n})) \right\} dk \\ &= \int_{k=0}^{\bar{n} \frac{f(\bar{n}) - h(\bar{n})}{f(1) - h(\bar{n})}} f(1) dk + \int_{k=\bar{n} \frac{f(\bar{n}) - h(\bar{n})}{f(1) - h(\bar{n})}}^{\bar{n}} \left(h(\bar{n}) + \frac{\bar{n}}{k} (f(\bar{n}) - h(\bar{n})) \right) dk \\ &= \bar{n} \frac{f(\bar{n}) - h(\bar{n})}{f(1) - h(\bar{n})} f(1) + \bar{n} \frac{f(1) - f(\bar{n})}{f(1) - h(\bar{n})} h(\bar{n}) + \bar{n} (f(\bar{n}) - h(\bar{n})) \ln \left(\frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})} \right) \end{aligned}$$

Dividing both sides of this inequality by $\bar{n}f(\bar{n})$ gives

$$\frac{F(\bar{n})}{\bar{n}f(\bar{n})} \leq 1 + \left(1 - \frac{h(\bar{n})}{f(\bar{n})} \right) \ln \left(\frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})} \right).$$

We want to determine what value for $f(\bar{n})$ maximizes the right-hand side. From the preceding, we know that $h(\bar{n}) \leq f(\bar{n}) \leq f(1)$ holds. The derivative with respect to the value of $f(\bar{n})$ is

$$\frac{d}{df(\bar{n})} \left(1 + \left(1 - \frac{h(\bar{n})}{f(\bar{n})} \right) \ln \left(\frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})} \right) \right) = \frac{h(\bar{n})}{f(\bar{n})^2} \ln \left(\frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})} \right) - \frac{1}{f(\bar{n})}.$$

For $f(\bar{n}) \searrow h(\bar{n})$ this goes to $+\infty$. For $f(\bar{n}) = f(1)$ this is $-\frac{1}{f(1)}$. Hence, the maximum is attained for $h(\bar{n}) < f(\bar{n}) < \bar{n}$. Setting the derivative to zero yields

$$\begin{aligned} \ln\left(\frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})}\right) &= \frac{f(\bar{n})}{h(\bar{n})} \Leftrightarrow \frac{f(1) - h(\bar{n})}{f(\bar{n}) - h(\bar{n})} = e^{\frac{f(\bar{n})}{h(\bar{n})}} \\ &\Leftrightarrow \frac{f(1) - h(\bar{n})}{eh(\bar{n})} = \frac{f(\bar{n}) - h(\bar{n})}{h(\bar{n})} e^{\frac{f(\bar{n}) - h(\bar{n})}{h(\bar{n})}} \\ &\Leftrightarrow f(\bar{n}) = h(\bar{n}) \left(1 + W_0\left(\frac{f(1) - h(\bar{n})}{eh(\bar{n})}\right)\right). \end{aligned}$$

As this is the unique point which satisfies the necessary conditions for a local optimum, by the preceding considerations it has to be indeed a global maximum. From substituting this value for $f(\bar{n})$, it follows

$$\frac{F(\bar{n})}{\bar{n}f(\bar{n})} \leq 1 + \frac{W_0\left(\frac{f(1) - h(\bar{n})}{eh(\bar{n})}\right)}{1 + W_0\left(\frac{f(1) - h(\bar{n})}{eh(\bar{n})}\right)} \ln\left(\frac{f(1) - h(\bar{n})}{h(\bar{n})W_0\left(\frac{f(1) - h(\bar{n})}{eh(\bar{n})}\right)}\right) = 1 + W_0\left(\frac{f(1) - h(\bar{n})}{eh(\bar{n})}\right).$$

As this bound is nondecreasing in \bar{n} , the maximum that appears in the bound of the potential function method is attained for $\bar{n} = n$. \blacksquare

The potential method for specific cost functions. Knowing the cost function allows for better bounds. We examine the potential method for affine and polynomial cost functions. The results are collected in [Figure 5.9](#).


In the case of affine total cost, we have $kf(k) = 1 + s(k - 1)$ for some $0 \leq s \leq 1$. Then, [Theorem 5.22](#) gives an upper bound on the price of stability of $1 - \ln((1 - s)/n + s)$. For $n \rightarrow +\infty$, this converges to a bound of $1 - \ln(s)$, which holds independently of the number of players. [Theorem 5.23](#) on the other hand gives $1 + W_0((1 - s)/(es))$. The inequalities for the principal branch of the Lambert W function by Hassani ([2005](#)) yield that $1 + W_0((1 - s)/(es)) \in \Theta(-\ln(s))$ for $s \rightarrow 0$. Hence, the asymptotic of the bounds from [Theorems 5.22](#) and [5.23](#) agree for $s \rightarrow 0$ and at $s = 1$. For fixed s , however, [Theorem 5.23](#) gives an improved estimate.

For the class of polynomial cost functions, [Theorems 5.22](#) and [5.23](#) give the bounds $1 + (1 - \alpha) \ln(n)$ and $1 + W_0((1 - n^\alpha + (n - 1)^\alpha)/(en^\alpha - e(n - 1)^\alpha))$ on the price of stability for games with n players. Again, carrying out the potential method more carefully allows to obtain the following constant bound.

Theorem 5.24 (The price of stability under polynomial cost). *Let $0 < \alpha \leq 1$. The price of stability of a network cost-sharing game with cost function $f(k) = k^{\alpha-1}$ is at most $1/\alpha$.*

Proof. For the specific cost function $f(k) = k^{\alpha-1}$, we get for $\bar{n} \in [n]$

$$\frac{F(\bar{n})}{\bar{n}f(\bar{n})} = \frac{\sum_{k=1}^{\bar{n}} k^{\alpha-1}}{\bar{n}^\alpha} \leq \frac{1 + \int_1^{\bar{n}} k^{\alpha-1} dk}{\bar{n}^\alpha} = \frac{1}{\alpha} - \frac{1 - \alpha}{\alpha \bar{n}^\alpha} \leq \frac{1}{\alpha}. \quad \blacksquare$$



(a) The two bounds $1 - \ln(s)$ (yellow) and $1 + W_0((1-s)/(es))$ (violet) on the price of stability under affine total cost. The bound of $1/s$ on the price of anarchy (green) is included for comparison.

(b) The bound $1/\alpha$ (red) on the price of stability under polynomial total edge cost. The bounds from Theorem 5.22 (yellow) and Theorem 5.23 (violet) are evaluated for $n = 100$. So is the price of anarchy (green).

Figure 5.9: Bounds on the price of stability under affine and polynomial total edge cost.

The potential method for the nonuniform case. The preceding results only deal with the uniform case of network cost-sharing games; that is, we assume that the cost function f is the same for all edges. In the nonuniform model, the cost function f_e is allowed to depend on e . The potential method relates the social cost and potential of a strategy profile for one edge at a time. Therefore, it easily generalizes to the nonuniform setting. The bound on the price of stability then is given by the worst bound on any cost function that appears. The following lemma gives some insight into the monotonicity of this bound.

Lemma 5.25 (Monotonicity of the potential function method). *Let $f, f': \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be two player cost functions such that $f'(k+1)/f'(1) - f'(k)/f'(1) \geq f(k+1)/f(1) - f(k)/f(1)$ for all $k \in \mathbb{N}$. Then, for all $n \in \mathbb{N}$*

$$\frac{\sum_{k=1}^n f'(k)}{nf'(n)} \leq \frac{\sum_{k=1}^n f(k)}{nf(n)}.$$

Proof. By induction, it follows that $f'(l)/f'(1) - f'(k)/f'(1) \geq f(l)/f(1) - f(k)/f(1)$ holds for all $1 \leq k \leq l$. In particular, it holds $f'(k)/f'(1) \geq f(k)/f(1)$ holds for all $k \in \mathbb{N}$. Applying these inequalities yields

$$\frac{\sum_{k=1}^n f'(k)}{nf'(n)} = 1 + \frac{\sum_{k=1}^n (f'(k) - f'(n))}{nf'(n)} \leq 1 + \frac{\sum_{k=1}^n (f(k) - f(n))}{nf(n)} = \frac{\sum_{k=1}^n f(k)}{nf(n)}. \quad \blacksquare$$

Lemma 5.25 suggests that the bound provided by the potential method does not depend so much on the absolute value of f than on its rate of change. This fits with the

observations for the extreme choices of the cost function. For $f(k) = 1$, which does not change at all and, thus, has the lowest possible rate of change, we obtain the best possible price of stability of 1. On the other side, $f(k) = 1/k$ decreases most rapidly (under our assumptions) and the potential method yields the worst possible bound of $H(n)$.

5.4.2.2 The Homogenization-Absorption Framework

The **homogenization-absorption framework** is inspired by the potential function method and applies to broadcast games with bounded total edge cost. It picks up the fundamental idea of starting with a social optimum σ^* and following improving steps to a Nash equilibrium. By choosing the improving steps carefully, it controls the quality of the resulting strategy profile. This is achieved by keeping it as close to the social optimum σ^* as possible. Two main concepts are involved in this process, **homogenization** and **absorption**. Ultimately, the homogenization-absorption framework is not an improving dynamics in the classical sense. The player cost of deviating players do not necessarily decrease. The improvement in each step rather lies in the strict decrease of the potential. In that sense, it is rather a guided version of the local search for the minimum concave-cost flow problem as discussed in [Section 5.3.1](#). Overall the process focuses on introducing as few edges from $E \setminus \text{supp}(\sigma^*)$ into the support of the strategy profile as possible. The underlying key observation is the following. Let T be the support of σ^* and S be the support of any other strategy profile $\sigma \in \Sigma$. For a cost function f with bounded $k \mapsto kf(k)$, the contribution of edges in T to the social cost of σ is only a constant factor away from the social optimum σ^* ; that is,

$$\sum_{e \in S \cap T} c_e n_\sigma(e) f(n_\sigma(e)) \leq \sup_{k \in [n]} \frac{kf(k)}{f(1)} \cdot \sum_{e \in S \cap T} c_e f(1) \leq \sup_{k \in \mathbb{N}} \frac{kf(k)}{f(1)} \cdot C(\sigma^*).$$

Hence, the ratio of the social costs of σ and σ^* is mostly determined by edges that σ uses but σ^* does not; that is,

$$\begin{aligned} \frac{C(\sigma)}{C(\sigma^*)} &= \frac{\sum_{e \in S \cap T} c_e n_\sigma(e) f(n_\sigma(e)) + \sum_{e \in S \setminus T} c_e n_\sigma(e) f(n_\sigma(e))}{C(\sigma^*)} \\ &\leq \sup_{k \in \mathbb{N}} \frac{kf(k)}{f(1)} + \frac{\sum_{e \in S \setminus T} c_e n_\sigma(e) f(n_\sigma(e))}{\sum_{e \in S \cap T} c_e n_{\sigma^*}(e) f(n_{\sigma^*}(e))} \\ &\leq \sup_{k \in \mathbb{N}} \frac{kf(k)}{f(1)} \cdot \left(1 + \frac{c(S \setminus T)}{c(T)} \right). \end{aligned}$$

In that sense, it is enough to control the use of edges outside of $\text{supp}(\sigma^*)$ and allow arbitrary usage of edges within $\text{supp}(\sigma^*)$. Even more, we can use edges in $\text{supp}(\sigma^*)$ in order to reduce the usage of edges in $E \setminus \text{supp}(\sigma^*)$.

The formal description of the algorithm is given in [Algorithm 5](#). The missing definitions are developed in this section. Note that we are not interested in the running time of the

Algorithm 5: The homogenization-absorption framework

Input: $G = (V, E), r \in V, c \in \mathbb{R}_{\geq 0}^E$, player cost function f
Output: Nash equilibrium σ

```

1 Compute a social optimum  $\sigma$ ;
2  $T \leftarrow \text{supp}(\sigma)$ ;
3 while  $\exists e = \{u, v\} \in E \setminus \text{supp}(\sigma) : c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) > c_e f(1)$  do
    // Absorb
4    $W \leftarrow \left\{ w \in V \mid d_{T, \bar{f}}(v, w) \leq \frac{f(1) - f(2)}{2} c_e \right\}$ ;
5    $\sigma \leftarrow \sigma^{W-T, e \rightarrow u}$ ;
    // Homogenize with respect to  $T$ 
6   while  $\exists w, z \in V : \Phi(\sigma^{w-T \rightarrow z}) < \Phi(\sigma)$  do
7      $\sigma \leftarrow \sigma^{w-T \rightarrow z}$ ;
8   end
9 end
10 return  $\sigma$ ;

```

algorithm but merely in the social cost of the resulting strategy profile. It starts with a social optimum and orchestrates deviations that utilize the support T of this initial strategy profile as much as possible. The main loop is executed as long as there is a deviating player v in the maintained profile σ . The stopping criterion is based on the characterization in [Corollary 5.6](#) for broadcast games. In the absorption step of every iteration, the neighborhood W of the deviating player v (with respect to a specific metric $d_{T, \bar{f}}$; see [Definition 5.28](#)) joins v on her new strategy via edges from T , yielding the strategy profile $\sigma^{W-T, e \rightarrow u}$ (see [Lemma 5.31](#)). This process ensures that within the neighborhood W the only edge in the support of the resulting strategy profile not in T is e , the edge certifying that the previous profile is not a Nash equilibrium. In order for the potential function to decrease during such an absorption, the players in W need to have similar total player costs to v in the previous profile. A sufficient condition thereof is the homogeneity of the profile (see [Definition 5.27](#)), which is sustained by an interleaved homogenization step. It uses deviations via edges in T that achieve a lower potential function value of the resulting strategy profile $\sigma^{w-T \rightarrow z}$ (see [Lemma 5.30](#)).

The following technical lemma is the main tool in the analysis of the strategy profiles resulting from homogenization and absorption. It relates the potential function value of a given strategy profile and the profile obtained after a cooperative deviation of a group of players. In the examined deviation, the strategies of the deviating players are changed to largely agree with the strategy of a specified player. For an illustration (using the notation of the lemma), see [Figure 5.10](#).




Figure 5.10: Illustration of the strategies of players $D_S(W)$ in the profiles σ and σ' from Lemma 5.26 with the set of deviating players W highlighted in yellow.

Lemma 5.26 (Change of potential under coordination). *Let $W \subseteq V$ and $v \in W$. Further, let $T \subseteq E$ be a tree such that $T[W]$ is connected. Let σ be a strategy profile supported by a spanning tree $S \subseteq E$ rooted at r , and define the strategy profile $\sigma' \in \Sigma$ for every $u \in V$ by*

$$\sigma'_u := \begin{cases} S[u, a_S^W(u)] \cup T[a_S^W(u), v] \cup S[v, r] & \text{if } u \in D_S(W), \\ S[u, r] & \text{otherwise.} \end{cases}$$

Then

$$\Phi(\sigma') - \Phi(\sigma) \leq \sum_{w \in W} |D_S^W(w)| (c'_\sigma(\sigma_v \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_v)) + \sum_{e \in T \setminus \sigma_v} c_e F\left(n_{\sigma'}^{D_S(W)}(e)\right)$$

where $n_{\sigma'}^{D_S(W)}(e) := |\{u \in D_S(W) \mid e \in \sigma'_u\}|$ denotes the number of players using an edge e in σ' when accounting only for descendants of W in S . Moreover, there is a strategy profile $\sigma'' \in \Sigma$ supported by a tree with $\Phi(\sigma'') \leq \Phi(\sigma')$.

Before giving the proof of Lemma 5.26, we remark that the tree T is generally chosen as a social optimum. Then the lemma is used to show that the potential decreases when a set of vertices W and their descendants cooperatively follow the social optimum to the strategy of a player v .

Proof. We start by estimating the difference of the potential for single edges. Thereby, we distinguish between three sets of edges. All deviating players join v on the edges σ_v . To do so, they use auxiliary edges from $T \setminus \sigma_v$. All other edges $S \setminus (\sigma_v \cup T)$ are abandoned by the cooperative deviation.

Joint edges. Every player u such that $\sigma_u \neq \sigma'_u$ uses all edges of σ_v in σ' , that is, $\sigma_v \subseteq \sigma'_u$. Therefore, we know $n_{\sigma'}(e) \geq n_{\sigma}(e)$ for all $e \in \sigma_v$. For all $e \in \sigma_v$, we use that f is nonincreasing to get

$$F(n_{\sigma'}(e)) - F(n_{\sigma}(e)) = \sum_{k=n_{\sigma}(e)+1}^{n_{\sigma'}(e)} f(k) \leq (n_{\sigma'}(e) - n_{\sigma}(e))f(n_{\sigma}(e) + 1).$$

Further, grouping the deviating players by their lowest ancestor in W yields

$$n_{\sigma'}(e) - n_{\sigma}(e) = |\{u \in D_S(W) \mid e \in \sigma'_u \setminus \sigma_u\}| = \sum_{\substack{w \in W: \\ e \in \sigma'_w \setminus \sigma_w}} |D_S^W(w)|.$$

Abandoned edges. An edge $e \in S \setminus (\sigma_v \cup T)$ is used in σ' at most as many times as in σ , that is, $n_{\sigma'}(e) \leq n_{\sigma}(e)$. This is due to the fact that $\sigma'_u \setminus \sigma_u \subseteq \sigma_v \cup T$ by construction. As f is nonincreasing, it holds for $e \in S \setminus (\sigma_v \cup T)$ that

$$F(n_{\sigma'}(e)) - F(n_{\sigma}(e)) = - \sum_{k=n_{\sigma'}(e)+1}^{n_{\sigma}(e)} f(k) \leq -(n_{\sigma}(e) - n_{\sigma'}(e))f(n_{\sigma}(e)).$$

For $w \in W$ and $u \in D_S^W(w)$, the definition of σ is such that $\sigma_w \setminus \sigma'_w = \sigma_u \setminus \sigma'_u$ holds. Further, $e \notin \sigma'_w \subseteq \sigma_v \cup T$. Therefore, we get that

$$n_{\sigma}(e) - n_{\sigma'}(e) = |\{u \in D_S(W) \mid e \in \sigma_u \setminus \sigma'_u\}| = \sum_{\substack{w \in W: \\ e \in \sigma_w}} |D_S^W(w)|.$$

Auxiliary edges. For an edge $e \in T \setminus \sigma_v$, there is no general monotonicity between $n_{\sigma'}(e)$ and $n_{\sigma}(e)$. Here it helps to compare σ and σ' to what they have in common, which is the strategies of players in $V \setminus D_S(W)$. Set $\hat{n} := |\{u \in V \setminus D_S(W) \mid e \in \sigma_u\}| = |\{u \in V \setminus D_S(W) \mid e \in \sigma'_u\}|$ to be the number of such players using e . Using again that f is nonincreasing yields for $e \in T \setminus \sigma_v$

$$F(n_{\sigma'}(e)) - F(n_{\sigma}(e)) = \sum_{k=\hat{n}+1}^{n_{\sigma'}(e)} f(k) - \sum_{k=\hat{n}+1}^{n_{\sigma}(e)} f(k) \leq F(n_{\sigma'}^{D_S(W)}(e)) - (n_{\sigma}(e) - \hat{n})f(n_{\sigma}(e)).$$

For $w \in W$ and $u \in D_S^W(w)$, we know $\sigma_w \subseteq \sigma_u$ by definition. Thus, it follows

$$n_\sigma(e) - \hat{n} = |\{u \in D_S(W) \mid e \in \sigma_u\}| \geq \sum_{\substack{w \in W: \\ e \in \sigma_w}} |D_S^W(w)|.$$

The change of potential. Finally, we put together the above estimates in order to get an upper bound on the difference in the potential. Note that the supports of σ and σ' lie in $S \cup T$. Thus, we get

$$\begin{aligned} & \Phi(\sigma') - \Phi(\sigma) - \sum_{e \in T \setminus \sigma_v} c_e F\left(n_{\sigma'}^{D_S(W)}(e)\right) \\ &= \sum_{e \in S \cup T} c_e \left(F(n_{\sigma'}(e)) - F(n_\sigma(e)) \right) - \sum_{e \in T \setminus \sigma_v} c_e F\left(n_{\sigma'}^{D_S(W)}(e)\right) \\ &\leq \sum_{\substack{e \in \sigma_v, \\ w \in W: \\ e \in \sigma'_w \setminus \sigma_w}} |D_S^W(w)| \cdot c_e f(n_\sigma(e) + 1) - \sum_{\substack{e \in (S \cup T) \setminus \sigma_v, \\ w \in W: \\ e \in \sigma_w}} |D_S^W(w)| \cdot c_e f(n_\sigma(e)) \\ &= \sum_{\substack{e \in \sigma_v, \\ w \in W: \\ e \in \sigma'_w \setminus \sigma_w}} |D_S^W(w)| \cdot c'_\sigma(e) - \sum_{\substack{e \in (S \cup T) \setminus \sigma_v, \\ w \in W: \\ e \in \sigma_w}} |D_S^W(w)| \cdot c_\sigma(e) \\ &= \sum_{w \in W} |D_S^W(w)| \cdot c'_\sigma(\sigma_v \setminus \sigma_w) - \sum_{w \in W} |D_S^W(w)| \cdot c_\sigma(\sigma_w \setminus \sigma_v). \end{aligned}$$

The last equality is based on changing the order of summation. This proves the first part of the lemma.

Reducing the support. The definition of σ' does not guarantee that the players' strategies are paths. By reducing them as follows, we claim to obtain a strategy profile σ'' that is supported by a tree. For all $u \in V \setminus D_S(W)$, set $\sigma''_u := \sigma'_u$. For every $u \in D_S(W)$, we choose a path $\sigma''_u \subseteq \sigma'_u$ such that $|\sigma''_u \setminus \sigma_v|$ is minimal. This definition directly implies $n_{\sigma''} \leq n_{\sigma'}$. Consequently, $\Phi(\sigma'') \leq \Phi(\sigma')$ as required. In order to prove the second part of the lemma, it remains to show that the support of σ'' indeed is a tree.

We first show that $\text{supp}(\sigma')$ is close to a tree. By the definition of $D_S(W)$, for all $u \in V \setminus D_S(W)$, the strategy $\sigma''_u = \sigma_u$ is not incident to any vertex in W . Similarly for $w \in W$, the tree $S[D_S^W(w)]$ is incident to no other vertex from W other than w . As $T[W]$ is a tree, we obtain that

$$T' := \bigcup_{u \in D_S(W)} (S[u, a_S^W(u)] \cup T[a_S^W(u), v]) = \bigcup_{\substack{w \in W, \\ u \in D_S^W(w)}} (S[u, w] \cup T[w, v]) = T[W] \cup \bigcup_{w \in W} S[D_S^W(w)]$$

is a tree again, which is vertex-disjoint to the strategies σ'_u of players $u \in V \setminus D_S(W)$. Hence, the support of σ' is the union of two vertex-disjoint trees and the path $S[v, r]$; that is,

$$\text{supp}(\sigma') = S[V \setminus D_S(W)] \cup T' \cup S[v, r].$$

By the preceding, σ'_u for $u \in D_S(W)$ is the union of the paths $S[u, a_S^W(u)] \cup T[a_S^W(u), v]$ and $S[v, r]$. Then the definition of σ''_u for such u yields that $\sigma''_u \setminus \sigma_v$ is a prefix of $\sigma'_u \setminus \sigma_v$ ending at the first vertex on σ_v ; that is,

$$\sigma''_u = T' \left[u, a_{T'}^{S[v,r]}(u) \right] \cup S \left[a_{T'}^{S[v,r]}(u), r \right]$$

where T' is considered to be rooted at v . Arguing analogously to T' being a tree, we get that $\bigcup_{u \in D_S[W]} \sigma''_u$ is a tree. On the other hand, $\sigma''_v \cup \bigcup_{u \in V \setminus D_S[W]} \sigma''_u$ is a subtree of S . Both, these trees make up $\text{supp}(\sigma'')$ and overlap only in σ_v . Hence, $\text{supp}(\sigma'')$ is indeed a tree. ■

Recall that the homogenization-absorption framework tries to maintain a strategy profile which is as close as possible to a social optimum. The maintained strategy profile is kept homogeneous with respect to a social optimum in the sense of the following definition.

Definition 5.27 (Homogeneity). *Let $\sigma \in \Sigma$ be a strategy profile that is supported by a spanning tree $S \subseteq E$. Further, let $T \subseteq E$ be another spanning tree. Define the strategy profile $\sigma^{w-T \rightarrow v} \in \Sigma$ by setting $\sigma_u^{w-T \rightarrow v} := \sigma_u$ for all $u \notin D_S(T[w, v])$ and choosing*

$$\sigma_u^{w-T \rightarrow v} \subseteq S \left[u, a_S^{T[w,v]}(u) \right] \cup T \left[a_S^{T[w,v]}(u), v \right] \cup S[v, r]$$

*to be a u - r path that minimizes $|\sigma_u^{w-T \rightarrow v} \setminus \sigma_v|$ for all $u \in D_S(T[w, v])$. The profile σ is called **homogeneous** with respect to T if $\Phi(\sigma) \leq \Phi(\sigma^{w-T \rightarrow v})$ holds for all $v, w \in V$.*

In other words, a homogeneous strategy σ with respect to a spanning tree $T \subseteq E$ is a local optimum with respect to the potential function and the neighborhood

$$N^{\text{homogeneous}}(\sigma) = \{ \sigma^{w-T \rightarrow v} \in \Sigma \mid w, v \in V \}.$$

The strategy profile $\sigma^{w-T \rightarrow v}$ results from σ when all descendants of vertices on the path from v to w in T deviate. Every such descendant u connects to this path $T[w, v]$ via a subpath of her strategy σ_u , then follows $T[w, v]$ to v , and finally follows the strategy of v in σ to r . The support of $\sigma^{w-T \rightarrow v}$ is reduced to a tree according to the last part of the proof of [Lemma 5.26](#). In a strategy profile that is homogeneous with respect to T , players that are close within T have similar total player costs. To express this proximity in T , we need the following notion of a distance.

Definition 5.28 (Broadcast distance). Let T be a spanning tree and $g: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be a nonincreasing function. We define the **directional broadcast distance** from vertex v to vertex w in T with respect to g as $\vec{d}_{T,g}(v, w) := \sum_{i=1}^k g(i)c_{e_i}$ where e_1, \dots, e_k are the edges on $T[v, w]$ in the order from v to w . Further, we define the **(undirectional) broadcast distance** between vertices v, w in T with respect to g as $d_{T,g}(v, w) := \sum_{i=1}^k g(i)c_{e_i}$ where e_1, \dots, e_k are the edges on $T[v, w]$ in descending order by cost, that is, $c_{e_1} \geq c_{e_2} \geq \dots \geq c_{e_k}$.

We use broadcast distances with respect to two functions. The first is the cost function f . For vertices v, w in a tree T , the directional broadcast distance $\vec{d}_{T,f}(v, w)$ gives an upper bound on the cost for v to travel to w in T while assuming that all players on the way follow along. The latter is a natural assumption as we know that a Nash equilibrium is supported by a tree. The second function with respect to which we measure the broadcast distance is $\bar{f}: \mathbb{N} \rightarrow \mathbb{R}_{>0}, k \mapsto F(k)/k$. So, $\bar{f}(k)$ is the average of the first k values of f . Since f is nonincreasing, also \bar{f} is nonincreasing. The directional broadcast distance arises naturally from the concept of homogeneity. The next lemma shows that the undirectional broadcast distance is metric, as opposed to the directional broadcast distance. This is why we use the undirectional version as an upper bound on the directional one.

Lemma 5.29 (Undirectional broadcast distance). Let $T \subseteq E$ be a spanning tree and $g: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be a nonincreasing function. For all $u, v, w \in V(T)$, it holds

$$(i) \quad d_{T,g}(u, v) \geq 0$$

$$(ii) \quad d_{T,g}(u, v) = d_{T,g}(v, u)$$

$$(iii) \quad d_{T,g}(u, w) \leq d_{T,g}(u, v) + d_{T,g}(v, w)$$

$$(iv) \quad d_{T,g}(u, v) \leq d_{T,g}(u, w) \text{ if } v \in V(T[u, w])$$

$$(v) \quad \vec{d}_{T,g}(u, v) \leq d_{T,g}(u, v) \leq d_{T,\bar{g}}(u, v) \text{ with } \bar{g}(k) := \frac{1}{k} \sum_{i=1}^k g(i) \text{ for } k \in \mathbb{N}_{>0}$$

Proof. Nonnegativity (i) follows directly from the definition of $d_{T,g}$. So does symmetry (ii), as $d_{T,g}$ is independent of the order of the edges between two vertices.

For a set of edges $P \subseteq E$, let $\pi_P: [|P|] \rightarrow P$ be a permutation of the edges in P in nonincreasing order by cost, that is, $c_{\pi_P(i)} \geq c_{\pi_P(j)}$ for all $1 \leq i < j \leq |P|$. This allows us to express the broadcast distance between two vertices u, v as

$$d_{T,g}(u, v) = \sum_{i=1}^{|T[u,v]|} g(i)c_{\pi_{T[u,v]}(i)} = \sum_{e \in T[u,v]} g\left(\pi_{T[u,v]}^{-1}(e)\right) c_e.$$

Monotonicity. To prove (iv), let $u, v, w \in V(T)$ such that v is on the path $T[u, w]$. From $T[u, v] \subseteq T[u, w]$, it follows that the $|T[u, v]|$ most expensive edges on $T[u, w]$

are at least as expensive as the edges on $T[u, v]$; that is, $c_{\pi_{T[u, v]}(i)} \leq c_{\pi_{T[u, w]}(i)}$ for all $i \in [T[u, v]]$. We use the preceding representation of $d_{T, g}$ to find

$$d_{T, g}(u, v) = \sum_{i=1}^{|T[u, v]|} g(i) c_{\pi_{T[u, v]}(i)} \leq \sum_{i=1}^{|T[u, v]|} g(i) c_{\pi_{T[u, w]}(i)} \leq d_{T, g}(u, w).$$

Triangle inequality. To see that (iii) holds, we fix three vertices $u, v, w \in V(T)$. As T is a tree, we know $T[u, w] \subseteq T[u, v, w] := T[u, v] \cup T[v, w]$ and, hence, we get by similar arguments as for monotonicity that

$$d_{T, g}(u, w) = \sum_{i=1}^{|T[u, w]|} g(i) c_{\pi_{T[u, w]}(i)} \leq \sum_{i=1}^{|T[u, v, w]|} g(i) c_{\pi_{T[u, v, w]}(i)} \leq \sum_{e \in T[u, v, w]} g\left(\pi_{T[u, v, w]}^{-1}(e)\right) c_e.$$

As g is nonincreasing, we continue

$$d_{T, g}(u, w) \leq \sum_{e \in T[u, v]} g\left(\pi_{T[u, v]}^{-1}(e)\right) c_e + \sum_{e \in T[v, w]} g\left(\pi_{T[v, w]}^{-1}(e)\right) c_e = d_{T, g}(u, v) + d_{T, g}(v, w).$$

This shows that $d_{T, g}$ indeed fulfills the triangle inequality.

Bounds. The first inequality of (v) is an immediate consequence of the rearrangement inequality as shown by Hardy, Littlewood, and Pólya (1953). The second inequality holds since g is nonincreasing and, therefore, $g \leq \bar{g}$. ■

Lemma 5.30 (Player costs under homogeneity). *Let $T \subseteq E$ be a spanning tree. Further, let $\sigma \in \Sigma$ be a strategy profile that is supported by a spanning tree $S \subseteq E$ and homogeneous with respect to T . Then for every $v, w \in V$, it holds*

$$c_{\sigma}(\sigma_v \setminus \sigma_w) - c'_{\sigma}(\sigma_w \setminus \sigma_v) \leq \vec{d}_{T, \bar{f}}(v, w).$$

Proof. The stated inequalities are a consequence of the local optimality of σ with respect to the neighborhood as defined by homogeneity. To see this, we set up a linear program with inequalities that result from applying Lemma 5.26 to σ and $\sigma^{w-T \rightarrow z}$ for all choices of $z \in V(T[w, v])$. Applying weak duality to the feasibility of the player costs in σ yields the statement.

Fix $v, w \in V$. Let the players and edges along the path $T[w, v]$ be labeled by $w = v_0, v_1, \dots, v_k = v$ and e_1, \dots, e_k , respectively, such that $e_i = \{v_{i-1}, v_i\}$ for all $i \in [k]$. Set $W := \{v_0, \dots, v_k\}$ to be the vertex set of $T[w, v]$. Further, let

$$n_j^{(i)} := \left| D_S^{T[v_0, v_i]}(v_j) \right| \quad \text{and} \quad N_j^{(i)} := \left| D_S^{T[v_0, v_i]}(T[v_0, v_j]) \right| = \sum_{k=0}^j n_k^{(i)}$$

for all $0 \leq j \leq i \leq k$.




Figure 5.11: The strategy $\sigma^{w-T \rightarrow v_i}$ as considered in the proof of Lemma 5.30. The descendants of $T[w, v_i]$ use the path $T[w, v]$ (yellow) to follow v_i (violet).

The linear program. By definition, the homogeneity of σ with respect to T implies $0 \leq \Phi(\sigma^{w-T \rightarrow v_i}) - \Phi(\sigma)$ for all $i \in \{0, \dots, k\}$. Note that $\sigma^{w-T \rightarrow v_i}$ can be obtained from σ by a construction as described in Lemma 5.26. See Figure 5.11 for an illustration of $\sigma^{w-T \rightarrow v_i}$. Hence, it follows

$$0 \leq \Phi(\sigma^{w-T \rightarrow v_i}) - \Phi(\sigma) \leq \sum_{j=0}^{i-1} n_j^{(i)} \left(c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \right) + \sum_{j=1}^i F \left(N_{j-1}^{(i)} \right) c_{e_j}.$$

Applying the estimate

$$c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \leq c_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) = c_\sigma(\sigma_{v_i}) - c_\sigma(\sigma_{v_j})$$

for all $i, j \in \{0, \dots, k\}$ shows that the choice $C_i = c_\sigma(\sigma_{v_i})$ for $i \in \{0, \dots, k\}$ is a feasible solution to the linear program

$$\begin{aligned} & \max \quad C_0 - C_k \\ & \text{s. t.} \quad \sum_{j=0}^{i-1} n_j^{(i)} (C_j - C_i) \leq \sum_{j=1}^i F \left(N_{j-1}^{(i)} \right) c_{e_j} \quad \text{for all } i \in [k] \\ & \quad \quad C_i \geq 0 \quad \quad \quad \text{for all } i \in \{0, \dots, k\}. \end{aligned}$$

In order to get an upper bound to the objective function value of any feasible solution, in particular to $c_\sigma(\sigma_{v_0}) - c_\sigma(\sigma_{v_k})$, we use weak duality. Dualizing the preceding linear

program gives

$$\begin{aligned}
 \min \quad & \sum_{j=1}^k \sum_{i=j}^k F(N_{j-1}^{(i)}) c_{e_j} y_i \\
 \text{s. t.} \quad & \sum_{j=1}^k n_0^{(j)} y_j \geq 1 \\
 & \sum_{j=i+1}^k n_i^{(j)} y_j - N_{i-1}^{(i)} y_i \geq 0 \quad \text{for all } i \in [k-1] \\
 & -N_{k-1}^{(k)} y_k \geq -1 \\
 & y_i \geq 0 \quad \text{for all } i \in [k].
 \end{aligned}$$

A feasible dual solution. Summing over all $k+1$ constraints of the dual (excluding the nonnegativity constraints) yields zero on both sides. Therefore, any feasible dual solution fulfills all of them with equality. Summing only the inequalities corresponding to C_i, \dots, C_k yields $\sum_{j=i}^k N_{i-1}^{(j)} y_j \leq 1$. Therefore, the recursive definition

$$\hat{y}_i := \left(N_{i-1}^{(i)} \right)^{-1} \left(1 - \sum_{j=i+1}^k N_{i-1}^{(j)} y_j \right)$$

for all $i \in [k]$ starting with $i = k$ gives the unique feasible dual solution.

A first bound. Weak duality implies that the primal objective value for $C_i = c_\sigma(\sigma_{v_i})$ is bounded from above by the dual objective value for $y_i = \hat{y}_i$. This gives

$$c_\sigma(\sigma_w \setminus \sigma_v) - c_\sigma(\sigma_v \setminus \sigma_w) = c_\sigma(\sigma_{v_0}) - c_\sigma(\sigma_{v_k}) \leq \sum_{j=1}^k \sum_{i=j}^k c_{e_j} F(N_{j-1}^{(i)}) \hat{y}_i$$

Because for fixed j the number $N_j^{(i)}$ is nonincreasing in i by its definition, and \bar{f} is nonincreasing as well, we obtain

$$\sum_{j=1}^k \sum_{i=j}^k c_{e_j} F(N_{j-1}^{(i)}) \hat{y}_i = \sum_{j=1}^k \sum_{i=j}^k \bar{f}(N_{j-1}^{(i)}) c_{e_j} N_{j-1}^{(i)} \hat{y}_i \leq \sum_{j=1}^k \bar{f}(N_{j-1}^{(k)}) c_{e_j} \sum_{i=j}^k N_{j-1}^{(i)} \hat{y}_i.$$

We have shown already that the inner sums of the right-hand side are at most one. We further use $N_{j-1}^{(k)} \geq j$ for all $j \in \mathbb{N}$ to find

$$\sum_{j=1}^k \bar{f}(N_{j-1}^{(k)}) c_{e_j} \sum_{i=j}^k N_{j-1}^{(i)} \hat{y}_i \leq \sum_{j=1}^k \bar{f}(N_{j-1}^{(k)}) c_{e_j} \leq \sum_{j=1}^k \bar{f}(j) c_{e_j} = \vec{d}_{T, \bar{f}}(w, v).$$

Note that combining these three chains of inequalities almost yields the inequality which we want to show. In order to improve it slightly, we construct another feasible solution to the primal linear program and apply the same reasoning to it.

Another feasible primal solution. We claim that $\widehat{C}_i := c'_\sigma(\sigma_{v_i} \cap \sigma_v) + c_\sigma(\sigma_{v_i} \setminus \sigma_v)$ for all $i \in \{0, \dots, k\}$ fulfills all inequalities of the primal linear program. We show $c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \leq \widehat{C}_i - \widehat{C}_j$ holds for all $i, j \in \{0, \dots, k\}$. Together with the inequalities we obtained from [Lemma 5.26](#), this proves the claim.

In the following, we leverage that S is a tree and, hence, the pairwise intersections of strategies follow some structure as indicated in [Figure 5.12](#). We start off by interpreting the relevant differences between strategies as part of the tree S through

$$\begin{aligned} & c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \\ &= c'_\sigma(S[v_i, \text{lca}(v_i, v_j)]) - c_\sigma(S[v_j, \text{lca}(v_i, v_j)]) \\ &= c'_\sigma(S[r, \text{lca}(v_i, v_j, v)]) + c_\sigma(S[\text{lca}(v_i, v_j, v), \text{lca}(v_i, v_j)]) + c'_\sigma(S[\text{lca}(v_i, v_j), v_i]) + \\ & \quad - c'_\sigma(S[r, \text{lca}(v_i, v_j, v)]) - c_\sigma(S[\text{lca}(v_i, v_j, v), \text{lca}(v_i, v_j)]) - c_\sigma(S[\text{lca}(v_i, v_j), v_j]). \end{aligned}$$

In the case of $\text{lca}(v_i, v_j, v) = \text{lca}(v_i, v_j)$, the right-hand side simplifies instantaneously and it follows that

$$\begin{aligned} & c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \\ &= c'_\sigma(S[r, v_i]) - c'_\sigma(S[r, \text{lca}(v_i, v_j, v)]) - c_\sigma(S[\text{lca}(v_i, v_j, v), v_j]) \\ &\leq c'_\sigma(S[r, \text{lca}(v_i, v)]) + c_\sigma(S[\text{lca}(v_i, v), v_i]) - c'_\sigma(S[r, \text{lca}(v_j, v)]) - c_\sigma(S[\text{lca}(v_j, v), v_j]) \\ &= \widehat{C}_i - \widehat{C}_j \end{aligned}$$

since $c'_\sigma \leq c_\sigma$ and $\text{lca}(v_i, v_j, v)$ is at least as close to the root r in S as $\text{lca}(v_j, v)$. Otherwise, $\text{lca}(v_i, v_j, v) \neq \text{lca}(v_i, v_j)$ implies $\text{lca}(v_i, v_j, v) = \text{lca}(v_i, v) = \text{lca}(v_j, v)$, as can be seen with the help of [Figure 5.12](#). Consequently, we obtain

$$\begin{aligned} & c'_\sigma(\sigma_{v_i} \setminus \sigma_{v_j}) - c_\sigma(\sigma_{v_j} \setminus \sigma_{v_i}) \\ &= c'_\sigma(S[r, \text{lca}(v_i, v)]) + c_\sigma(S[\text{lca}(v_i, v), \text{lca}(v_i, v_j)]) + c'_\sigma(S[\text{lca}(v_i, v_j), v_i]) + \\ & \quad - c'_\sigma(S[r, \text{lca}(v_j, v)]) - c_\sigma(S[\text{lca}(v_j, v), \text{lca}(v_i, v_j)]) - c_\sigma(S[\text{lca}(v_i, v_j), v_j]) \\ &\leq c'_\sigma(S[r, \text{lca}(v_i, v)]) + c_\sigma(S[\text{lca}(v_i, v), \text{lca}(v_i, v_j)]) + c_\sigma(S[\text{lca}(v_i, v_j), v_i]) + \\ & \quad - c'_\sigma(S[r, \text{lca}(v_j, v)]) - c_\sigma(S[\text{lca}(v_j, v), \text{lca}(v_i, v_j)]) - c_\sigma(S[\text{lca}(v_i, v_j), v_j]) \\ &= c'_\sigma(S[r, \text{lca}(v_i, v)]) + c_\sigma(S[\text{lca}(v_i, v), v_i]) - c'_\sigma(S[r, \text{lca}(v_j, v)]) - c_\sigma(S[\text{lca}(v_j, v), v_j]) \\ &= \widehat{C}_i - \widehat{C}_j. \end{aligned}$$




Figure 5.12: Possible relationships between the lowest common ancestors in S .

In total, we see that $\widehat{C}_i, i \in \{0, \dots, k\}$ indeed give a feasible solution to the linear program. Proceeding as before, we use the feasible dual solution $\widehat{y}_i, i \in [k]$ and weak duality to eventually find

$$c_\sigma(\sigma_w \setminus \sigma_v) - c'_\sigma(\sigma_v \setminus \sigma_w) = c'_\sigma(\sigma_w \cap \sigma_v) + c_\sigma(\sigma_w \setminus \sigma_v) - c'_\sigma(\sigma_v) = \widehat{C}_0 - \widehat{C}_k \leq \vec{d}_{T, \vec{f}}(w, v). \quad \blacksquare$$

Now we are able to analyze the absorption step. When a player deviates from a strategy profile that is homogeneous with respect to T , all players that are close enough with respect to the broadcast distance in T can follow while the overall change of the potential is negative.

Lemma 5.31 (Absorption). *Let $T \subseteq E$ be a spanning tree. Further, let $\sigma \in \Sigma$ be a strategy profile which is supported by a spanning tree $S \subseteq E$ and is homogeneous with respect to T . Assume that there is an edge $e = \{u, v\} \in E \setminus (S \cup T)$ such that $c'_\sigma(\sigma_u \setminus \sigma_v) + c_e f(1) < c_\sigma(\sigma_v \setminus \sigma_u)$. Define the set*

$$W := \left\{ w \in V \mid d_{T, \vec{f}}(v, w) \leq \frac{f(1) - f(2)}{2} c_e \right\}$$

and the strategy profile $\sigma^{W-T, e \rightarrow u}$ by setting for $z \in V$

$$\sigma_z^{W-T, e \rightarrow u} := \begin{cases} S[z, a_S^W(z)] \cup T[a_S^W(z), v] \cup \{e\} \cup S[u, r] & \text{if } z \in D_S(W), \\ S[z, r] & \text{otherwise.} \end{cases}$$

Then it holds $\Phi(\sigma^{W-T, e \rightarrow u}) < \Phi(\sigma)$. Moreover, for every $w \in W \setminus \{v\}$ with $e_S(w) \notin T$ the edge $e_S(w)$ is not in the support of $\sigma^{W-T, e \rightarrow u}$.

Proof. We combine [Lemmas 5.26](#) and [5.30](#) to show that the potential of $\sigma^{W-T, e \rightarrow u}$ is strictly less than the potential of σ .

Relating to u . First, we show that $\sigma^{W-T, e \rightarrow u}$ can be obtained from σ by a construction as in [Lemma 5.26](#). By [Lemma 5.29 \(iv\)](#), $T[W]$ is a tree. We show that $V(\sigma_u) \cap W = \emptyset$. To that end, let $z \in V(\sigma_u)$. Then $\sigma_z \subseteq \sigma_u$ shows $c_\sigma(\sigma_v \setminus \sigma_u) \leq c_\sigma(\sigma_v \setminus \sigma_z)$ and $c'_\sigma(\sigma_u \setminus \sigma_v) \geq c'_\sigma(\sigma_z \setminus \sigma_v)$. The assumption that v wants to deviate via e in combination




Figure 5.13: The contributions of paths in S to $c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u)$. Paths highlighted in yellow add to $c'_\sigma(\sigma_u \setminus \sigma_w)$. Paths highlighted in violet add to $c_\sigma(\sigma_w \setminus \sigma_u)$.

with Lemma 5.30 yields

$$\begin{aligned} c_e \frac{f(1) - f(2)}{2} &\leq c_e f(1) < c_\sigma(\sigma_v \setminus \sigma_u) - c'_\sigma(\sigma_u \setminus \sigma_v) \\ &\leq c_\sigma(\sigma_v \setminus \sigma_z) - c'_\sigma(\sigma_z \setminus \sigma_v) \leq \vec{d}_{T, \bar{f}}(u, v) \leq d_{T, \bar{f}}(u, v). \end{aligned}$$

Thus, $z \notin W$ as claimed. In particular, $T[W] \cup \{e\}$ is a tree again and Lemma 5.26 is applicable. The support of $\sigma^{W-T, e \rightarrow u}$ does not change by the reduction in the proof of Lemma 5.26, because $V(\sigma_u) \cap W = \emptyset$ holds. Hence, it is a tree. Particularly, all edges $e_S(w) \notin T$ for $w \in W \setminus \{v\}$ are removed from the support. Further, Lemma 5.26 implies

$$\begin{aligned} \Phi(\sigma^{W-T, e \rightarrow u}) - \Phi(\sigma) &\leq \sum_{w \in W} |D_S^W(w)| \left(c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u) \right) \\ &\quad + \sum_{e' \in T[W] \setminus \sigma_u} c_{e'} F \left(n_{\sigma^{W-T, e \rightarrow u}}^{D_S(W)}(e') \right) + c_e F \left(|D_S(W)| \right). \end{aligned}$$

Relating to v . The term on the right hand side contains the cost differences $c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u)$ while the definition of W in combination with Lemma 5.30 results in estimates for the cost differences $c'_\sigma(\sigma_v \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_v)$. Therefore, we relate these two differences to each other. For every $w \in W$, we use the fact that S is a tree and that v wants to deviate via e . Thus it holds

$$\begin{aligned} c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u) &= c'_\sigma(S[\text{lca}(u, w), u]) - c_\sigma(S[\text{lca}(u, w), w]) \\ &= c'_\sigma(S[\text{lca}(u, v, w), \text{lca}(u, v)]) + c'_\sigma(S[\text{lca}(u, v), u]) \\ &\quad - c'_\sigma(S[\text{lca}(u, v, w), \text{lca}(u, w)]) - c_\sigma(S[\text{lca}(u, w), w]) \\ &< c'_\sigma(S[\text{lca}(u, v, w), \text{lca}(u, v)]) + c_\sigma(S[\text{lca}(u, v), v]) \\ &\quad - c'_\sigma(S[\text{lca}(u, v, w), \text{lca}(u, w)]) - c_\sigma(S[\text{lca}(u, w), w]) - c_e f(1). \end{aligned}$$

If $\text{lca}(u, v, w) = \text{lca}(v, w)$, then it follows

$$\begin{aligned} c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u) &< c'_\sigma(S[\text{lca}(v, w), \text{lca}(u, v)]) + c_\sigma(S[\text{lca}(u, v), v]) \\ &\quad - c'_\sigma(S[\text{lca}(v, w), \text{lca}(u, w)]) - c_\sigma(S[\text{lca}(u, w), w]) - c_e f(1) \\ &\leq c_\sigma(S[\text{lca}(v, w), v]) - c'_\sigma(S[\text{lca}(v, w), w]) - c_e f(1) \\ &\leq c_\sigma(\sigma_v \setminus \sigma_w) - c'_\sigma(\sigma_w \setminus \sigma_v) - c_e f(1). \end{aligned}$$

Otherwise, we are in the case of [Figure 5.13a](#), that is, $\text{lca}(u, v, w) = \text{lca}(u, v) = \text{lca}(u, w)$ holds. Then we get

$$\begin{aligned} c'_\sigma(\sigma_u \setminus \sigma_w) - c_\sigma(\sigma_w \setminus \sigma_u) &< c_\sigma(S[\text{lca}(u, v, w), v]) - c_\sigma(S[\text{lca}(u, v, w), w]) - c_e f(1) \\ &= c_\sigma(S[\text{lca}(v, w), v]) - c_\sigma(S[\text{lca}(v, w), w]) - c_e f(1) \\ &\leq c_\sigma(S[\text{lca}(v, w), v]) - c'_\sigma(S[\text{lca}(v, w), w]) - c_e f(1) \\ &\leq c_\sigma(\sigma_v \setminus \sigma_w) - c'_\sigma(\sigma_w \setminus \sigma_v) - c_e f(1). \end{aligned}$$

So we obtain the same inequality in both cases. Plugging it into the upper bound on the difference of the potentials of σ and $\sigma^{W-T, e \rightarrow u}$ yields

$$\begin{aligned} \Phi(\sigma^{W-T, e \rightarrow u}) - \Phi(\sigma) &< \sum_{w \in W} |D_S^W(w)| \left(c_\sigma(\sigma_v \setminus \sigma_w) - c'_\sigma(\sigma_w \setminus \sigma_v) - c_e f(1) \right) \\ &\quad + \sum_{e' \in T[W] \setminus \sigma_u} c_{e'} F \left(n_{\sigma^{W-T, e \rightarrow u}}^{D_S(W)}(e') \right) + c_e F \left(|D_S(W)| \right). \end{aligned}$$

The inequality is strict as $v \in W$ and $v \in D_S^W(v) \neq \emptyset$.

Relating to the broadcast distance. The terms on the right-hand side of the preceding inequality can be bounded in terms of the broadcast distance. The definition of W then yields an upper bound of zero. As we assume that σ is homogeneous, we can apply [Lemma 5.30](#) to the first of the three summands on the right-hand side. In order to bound the remaining terms, we carefully analyze the contribution of players to the congestion of e and edges in $T[W]$. To that end, let $T' \subseteq E$ be the set of edges used by players of $D_S(W)$ to connect to v in $\sigma^{W-T, e \rightarrow u}$, that is, $T' = T[W] \cup \bigcup_{w \in W} S[D_S(w)]$. Similar to the last part of the proof of [Lemma 5.26](#), it can be shown that T' is a tree on the vertices $D_S(W)$ rooted at v . For any $0 \leq k < |T'|$, the rooted tree T' has at least $k + 1$ vertices with depth less than or equal to k . This is a simple consequence of connectedness. Using in addition that f is nonincreasing yields

$$F \left(|D_S(W)| \right) = \sum_{k=1}^{|D_S(W)|} f(k) \leq \sum_{z \in D_S(W)} f(|T'[z, v]| + 1) \leq \sum_{w \in W} |D_S^W(w)| f(|T[w, v]| + 1).$$

We denote the path in T' from a vertex $z \in D_S(W)$ up to (but without) some edge $e' \in T'$ by $T'[z, e']$. By reasoning for subtrees in T' underneath edges $e' \in T[W]$ just like for T' , we get

$$F\left(n_{\sigma^{W-T, e \rightarrow u}}^{D_S(W)}(e')\right) \leq \sum_{\substack{z \in D_S(W): \\ e' \in T'[z, v]}} f(|T'[z, e']| + 1) \leq \sum_{\substack{w \in W: \\ e' \in T[w, v]}} |D_S^W(w)| f(|T'[w, e']| + 1).$$

We combine these inequalities for all $e' \in T[W] \setminus \sigma_u$ to obtain

$$\begin{aligned} \sum_{e' \in T[W] \setminus \sigma_u} c_{e'} F\left(n_{\sigma^{W-T, e \rightarrow u}}^{D_S(W)}(e')\right) &\leq \sum_{\substack{e' \in T[W] \setminus \sigma_u, \\ w \in W: \\ e' \in T[w, v]}} |D_S^W(w)| \cdot c_{e'} f(|T'[w, e']| + 1) \\ &= \sum_{w \in W} |D_S^W(w)| \sum_{e' \in T[w, v]} c_{e'} f(|T[w, e']| + 1) \\ &= \sum_{w \in W} |D_S^W(w)| \vec{d}_{T, f}(w, v). \end{aligned}$$

The overall bound. Combining the inequalities of the preceding paragraph with our bound on the change of the potential results in

$$\begin{aligned} &\Phi(\sigma^{W-T, e \rightarrow u}) - \Phi(\sigma) \\ &< \sum_{w \in W} |D_S^W(w)| \left(\vec{d}_{T, \bar{f}}(v, w) - c_e f(1) + \vec{d}_{T, f}(w, v) + c_e f(|T[w, v]| + 1) \right) \\ &\leq \sum_{w \in W \setminus \{v\}} |D_S^W(w)| \left(\vec{d}_{T, \bar{f}}(v, w) + \vec{d}_{T, f}(w, v) + (f(2) - f(1))c_e \right). \end{aligned}$$

Finally, [Lemma 5.29 \(v\)](#) and [\(ii\)](#) together the definition of W based on the distance $d_{T, \bar{f}}$ imply that every single summand on the right-hand side is nonpositive and, therefore, also $\Phi(\sigma^{W-T, e \rightarrow u}) - \Phi(\sigma) < 0$. \blacksquare

Theorem 5.32 (Correctness and termination). *Algorithm 5 returns a Nash equilibrium after finitely many steps.*

Proof. To prove correctness of [Algorithm 5](#), we show that the support of σ throughout the execution of the algorithm is a spanning tree. As shown before, [Lemma 5.26](#) applies to the strategy profiles $\sigma^{w-T \rightarrow z}$ and $\sigma^{W-T, e \rightarrow u}$ resulting from homogenization and absorption, respectively. These profiles are supported by trees. Then, correctness of [Algorithm 5](#) is an immediate consequence of [Corollary 5.6](#) in combination with the termination criterion of the main loop.

The algorithm's progress is measured by the decrease of the potential. By [Lemma 5.31](#), the potential Φ of the current strategy profile decreases strictly in each absorption step. The same is true for each iteration of the homogenization. In total, homogenization

does not increase the potential. Hence, no strategy profile appears in more than one iteration. Because there are only finitely many strategy profiles, [Algorithm 5](#) terminates after finitely many iterations. ■

Theorem 5.33 (The price of stability under bounded cost). *Let f be strictly decreasing. For any $0 < \rho' < \rho < \frac{1}{2}$, the Nash equilibrium computed by [Algorithm 5](#) has social cost at most*

$$\left(\sup_{k \in \mathbb{N}} \frac{kf(k)}{f(1)} \right) \left(1 + \frac{2f(1)}{f(1) - f(2)} \left(\frac{(1 - \rho)\rho}{(1 - 2\rho)\rho'^2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 + \frac{1 - \rho'}{(1 - 2\rho')(\rho - \rho')} \right) \right)$$

times the cost of a social optimum.

Proof. Let $T \subseteq E$ be the minimum spanning tree computed in the beginning of the algorithm. Further, let σ be the Nash equilibrium computed by [Algorithm 5](#) and $S \subseteq E$ be its supporting spanning tree. We bound the social cost of σ with respect to the total cost of the edges in T , which is equal to $\sum_{e \in T} c_e$. The latter can be bounded with respect to the social cost of a social optimum, which then yields the statement of the theorem. To estimate the contribution $c_e n_\sigma(e) f(n_\sigma(e))$ of edges $e \in S \cap T$ to the social cost of σ , we use the trivial bound

$$c_e n_\sigma(e) f(n_\sigma(e)) \leq \sup_{k \in \mathbb{N}} \frac{kf(k)}{f(1)} \cdot c_e f(1).$$

In the remaining, we focus on bounding the contribution by edges in $S \setminus T$ with the help of a charging scheme. To that end, let the edges $S \setminus T = \{e_1, \dots, e_k\}$ be indexed in the order that they were added to S by [Algorithm 5](#). If an edge is added to S in more than one iteration (since it is removed in between), we consider the last point in time that it is being added. Note that neither the absorption nor the homogenization step introduce edges from $E \setminus T$ into the support of the strategy profile except for the respective e_i . Further, let v_1, \dots, v_k and u_1, \dots, u_k be the vertices such that $e_i = \{u_i, v_i\} = e_S(v_i)$ for all $i \in [k]$, that is, u_i is closer to the root r in S than v_i . Note that homogenization and absorption (in particular, the reduction of the support according to [Lemma 5.26](#)) preserve the direction in which an edge in $S \setminus T$ is used, that is, the definition of u_i and v_i .

The charging scheme. Let $\alpha = \frac{1}{2}(1 - f(2)/f(1))$ denote the relative absorption radius. We apply a charging scheme in order to bound $c(S \setminus T)$ by a multiple of $c(T)$; that is, we assign the edge costs in $S \setminus T$ to edges in T . We charge the contribution of e_i to the social cost to edges in $T[v_i, r]$ which are relatively close to v_i . This proximity is specified by $0 < \rho < \frac{1}{2}$. For $i \in [k]$, let $w_i \in T[v_i, r]$ such that $|T[v_i, w_i]|$ is maximal subject to $d_{T, \bar{f}}(v_i, w_i) \leq \rho \alpha c_{e_i} f(1)$. We charge $T[v_i, w_i]$ and the adjacent edge $e_T(w_i)$ for e_i . To see that $e_T(w_i)$ exists, we need to show that $w_i \neq r$. For that purpose, let $\sigma^{(i)}$ denote the strategy profile at the beginning of the iteration in which e_i was added to S for the last time. By the construction of the algorithm, $\sigma^{(i)}$ is homogeneous with respect to T . Note

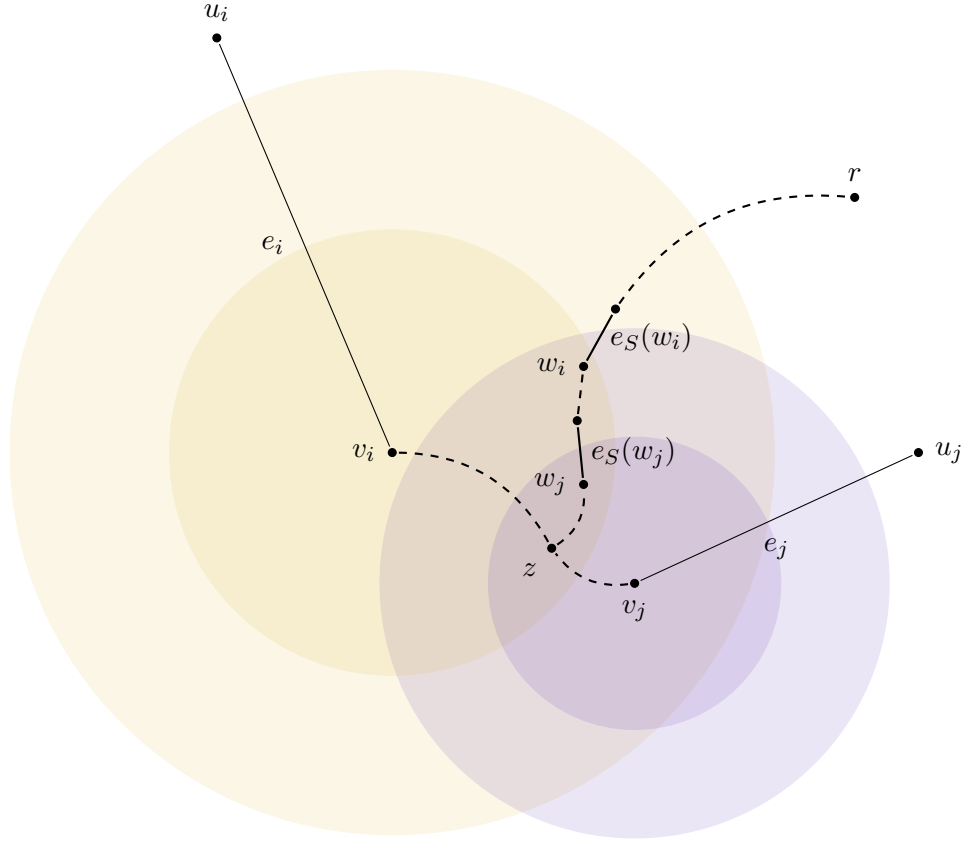


Figure 5.14: Charging scheme from the proof of [Theorem 5.33](#). The **yellow** circles represent the absorption radius $\alpha c_{e_i} f(1)$ and the charging radius $\rho \alpha c_{e_i} f(1)$ around v_i with respect to $d_{T, \bar{f}}$. Similarly, the **violet** circles represent the absorption radius $\alpha c_{e_j} f(1)$ and the charging radius $\rho \alpha c_{e_j} f(1)$ around v_j with respect to $d_{T, \bar{f}}$. The thick connections represent edges and paths in T . The vertex z is on both paths, $T[v_i, w_j]$ and $T[v_j, w_j]$.

that the profile induced by T is trivially homogeneous with respect to T . It holds by the choice of $e = e_i$ and [Lemma 5.30](#) that

$$c_{e_i} \leq c'_{\sigma^{(i)}}(\sigma_{u_i}^{(i)} \setminus \sigma_{v_i}^{(i)}) + c_{e_i} < c_{\sigma^{(i)}}(\sigma_{v_i}^{(i)} \setminus \sigma_{u_i}^{(i)}) \leq c_{\sigma^{(i)}}(\sigma_{v_i}^{(i)} \setminus \sigma_r^{(i)}) \leq \vec{d}_{T, \bar{f}}(v_i, r) \leq d_{T, \bar{f}}(v_i, r)$$

where $\sigma_r^{(i)}$ is considered to be the empty strategy. By the definition of w_i , it follows that $w_i \neq r$ and we can indeed charge $e_T(w_i)$. [Figure 5.14](#) provides a schematic overview of the charging scheme.

Now we define the exact value $\gamma_{e_i, e}$ that an edge $e \in T$ is charged for an edge $e_i \in S \setminus T, i \in [k]$. Fix $0 < \rho' < \rho$, which we use to decide whether to charge $T[v_i, w_i]$ or $e_T(w_i)$. In the case of $d_{T, \bar{f}}(v_i, w_i) \geq \rho' \alpha c_{e_i} f(1)$, the cost of e_i is distributed to the edges

of $T[v_i, w_i]$ by

$$\gamma_{e_i, e} := \begin{cases} (\rho'\alpha)^{-2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 \frac{c_e}{c_{e_i}} c_e & \text{if } e \in T[v_i, w_i], \\ 0 & \text{otherwise.} \end{cases}$$

Otherwise, the whole cost of e_i is charged to $e_T(w_i)$ by

$$\gamma_{e_i, e} := \begin{cases} \frac{c_{e_i}}{c_e} c_e & \text{if } e = e_T(w_i), \\ 0 & \text{otherwise.} \end{cases}$$

Cost coverage. We first show that the cost of an edge e_i for $i \in [k]$ is covered by the charges. If $d_{T, \bar{f}}(v_i, w_i) < \rho'\alpha c_{e_i} f(1)$, this is clearly the case because $\sum_{e \in T} \gamma_{e_i, e} = c_{e_i}$. Otherwise, squaring the inequality $d_{T, \bar{f}}(v_i, w_i) \geq \rho'\alpha c_{e_i} f(1)$ and applying the Cauchy-Schwarz inequality to $d_{T, \bar{f}}(v_i, w_i)^2$ yields

$$c_{e_i} \leq (\rho'\alpha)^{-2} \frac{d_{T, \bar{f}}(v_i, w_i)^2}{c_{e_i} f(1)^2} \leq (\rho'\alpha)^{-2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 \sum_{e \in T[v_i, w_i]} \frac{c_e^2}{c_{e_i}} = \sum_{e \in T} \gamma_{e_i, e}.$$

Overcharge. As the paths $T[v_i, w_i] \cup \{e_T(w_i)\}$ can generally not be expected to be pairwise disjoint, we examine which e_i are charged to the same edge in T . For $i < j$ such that there is $z \in V(T[v_i, w_i]) \cap V(T[v_j, w_j])$, the triangle inequality (Lemma 5.29 (iii)) and the monotonicity property of $d_{T, \bar{f}}$ (Lemma 5.29 (iv)) yield

$$d_{T, \bar{f}}(v_i, v_j) \leq d_{T, \bar{f}}(v_i, z) + d_{T, \bar{f}}(v_j, z) \leq d_{T, \bar{f}}(v_i, w_i) + d_{T, \bar{f}}(v_j, w_j) \leq \rho\alpha(c_{e_i} + c_{e_j})f(1).$$

On the other hand, we know that e_i was not absorbed in the iteration when e_j was added to S for the last time. Thus, we know $d_{T, \bar{f}}(v_i, v_j) > \alpha c_{e_j} f(1)$. Combined with the preceding inequality, this gives $c_{e_j} < \rho/(1 - \rho) \cdot c_{e_i}$ for all $1 \leq i < j \leq k$ with $V(T[v_i, w_i]) \cap V(T[v_j, w_j]) \neq \emptyset$.

If $e_T(w_i) = e_T(w_j)$ for $1 \leq i < j \leq k$ then $w_i = w_j$. If, in addition, both e_i and e_j get charged to $e_T(w_i) = e_T(w_j)$, then $d_{T, \bar{f}}(v_i, w_i) < \rho'\alpha c_{e_i} f(1)$ and $d_{T, \bar{f}}(v_j, w_j) < \rho'\alpha c_{e_j} f(1)$. Arguing as in the preceding paragraph, it follows that $c_{e_j} < \rho'/(1 - \rho') \cdot c_{e_i}$ in this case.

Next we bound the ratios of the costs of charging and charged edge. For $i \in [k]$ and $e \in T[v_i, w_i]$, we get

$$c_e f(1) \leq \max_{e' \in T[v_i, w_i]} c_{e'} f(1) \leq d_{T, \bar{f}}(v_i, w_i) \leq \rho\alpha c_{e_i} f(1).$$

On the other hand, let $i \in [k]$ such that $d_{T, \bar{f}}(v_i, w_i) < \rho'\alpha c_{e_i} f(1)$ and $e = e_T(w_i)$. Due to the choice of w_i and the triangle inequality, we obtain

$$\rho\alpha c_{e_i} f(1) < d_{T, \bar{f}}(v_i, p_T(w_i)) \leq d_{T, \bar{f}}(v_i, w_i) + c_e f(1) < (\rho'\alpha c_{e_i} + c_e) f(1)$$

where $p_T(w_i)$ denotes the parent of w_i , that is, its adjacent ancestor in the rooted tree T . Rearranging yields $(\rho - \rho')\alpha c_{e_i} < c_e$.

Now, we can bound the overcharge of edges in T under this scheme. For $e \in T$, the charging of the two types yields

$$\sum_{i=1}^k \gamma_{e_i, e} = (\rho' \alpha)^{-2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 \sum_{\substack{i \in [k]: \\ e \in T[v_i, w_i]}} \frac{c_e}{c_{e_i}} c_e + \sum_{\substack{i \in [k]: \\ e = e_T(w_i)}} \frac{c_{e_i}}{c_e} c_e$$

Applying the bounds on the ratios between charging and charged edges, and the bounds on the ratio between different charging edges gives

$$\begin{aligned} \sum_{i=1}^k \gamma_{e_i, e} &\leq \frac{\rho}{\rho'^2 \alpha} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 \sum_{l=0}^{\infty} \left(\frac{\rho}{1-\rho} \right)^l c_e + \frac{1}{(\rho - \rho') \alpha} \sum_{l=0}^{\infty} \left(\frac{\rho'}{1-\rho'} \right)^l c_e \\ &\leq \left(\frac{(1-\rho)\rho}{(1-2\rho)\rho'^2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 + \frac{1-\rho'}{(1-2\rho')(\rho - \rho')} \right) \alpha^{-1} c_e. \end{aligned}$$

The overall bound. In total, we get

$$\begin{aligned} c(S) &\leq c(T) + c(S \setminus T) \leq c(T) + \sum_{e \in T} \sum_{i=1}^k \gamma_{e_i, e} \\ &\leq \left(1 + \alpha^{-1} \left(\frac{(1-\rho)\rho}{(1-2\rho)\rho'^2} \sum_{j=1}^{\infty} \left(\frac{\bar{f}(j)}{f(1)} \right)^2 + \frac{1-\rho'}{(1-2\rho')(\rho - \rho')} \right) \right) c(T) \end{aligned}$$

Observing $C(S) \leq \sup_{l \in \mathbb{N}} l f(l) c(S)$ and $f(1) c(T) \leq C(T)$ finishes the proof. \blacksquare

Corollary 5.34 (The price of stability under constant total edge cost). *For the cost function $f(k) = 1/k$, the price of stability for broadcast games is at most 265.*

Proof. For $f(k) = \frac{1}{k}$, we have $\sup_{l \in \mathbb{N}} l f(l) / f(1) = 1$. Further, the series $\sum_{j=1}^{\infty} \bar{f}(j)^2 = \sum_{j=1}^{\infty} H(j)^2 / j^2$ belongs to the so-called Euler sums. It evaluates to $\frac{17}{4} \zeta(4) = \frac{17\pi^4}{360}$ as shown by D. Borwein and J. M. Borwein (1995) (where ζ denotes the Riemann zeta function). Hence, the upper bound from [Theorem 5.33](#) evaluates to

$$1 + 4 \left(\frac{(1-\rho)\rho}{(1-2\rho)\rho'^2} \frac{17\pi^4}{360} + \frac{1-\rho'}{(1-2\rho')(\rho - \rho')} \right)$$

for any $0 < \rho' < \rho < \frac{1}{2}$. Setting $\rho = 0.31$, $\rho' = 0.24$ yields an upper bound below 265. \blacksquare

5.4.3 Lower Bounds on the Price of Stability

In this section, we focus on the price of stability of broadcast games. It is defined as the maximal price of stability of any instance. The price of stability of an instance, in turn, equals the ratio of the minimal cost of any Nash equilibrium and the cost of a social optimum. Therefore, we are facing the max-min problem

$$\text{PoS} = \max_{\substack{G=(V,E) \\ \text{graph}}} \max_{c \in \mathbb{R}_{\geq 0}^E} \min_{\substack{\sigma \in \Sigma \\ \text{Nash equilibrium}}} \frac{C(\sigma)}{C(\sigma^*)}$$

where σ^* denotes a social optimum in G with respect to the cost c . In the following, we successively restrict the structure of the graph G and the cost c in order to simplify the problem while maintaining a lower bound. Dealing with the internal minimization problem of finding a minimum-cost Nash equilibrium seems the most challenging. We take an approach similar to the one of Bilò, Caragiannis, et al. (2013). We restrict the graph and the edge costs in such a way that the inner problem becomes trivial. This can be achieved by designating a minimum-cost Nash equilibrium first and then choosing the cost which complies with this designation. In some sense, this separates the structural properties from the numerical ones. We discuss the price of stability under various fixed structural features. The simplest one considering a fixed graph only. We denote the price of stability on a graph G by

$$\text{PoS}(G) = \max_{c \in \mathbb{R}_{\geq 0}^E} \min_{\substack{\sigma \in \Sigma \\ \text{Nash equilibrium}}} \frac{C(\sigma)}{C(\sigma^*)}.$$

Further, fixing the support $S \subseteq E$ of the designated minimum-cost Nash equilibrium yields the price of stability on G conditioning on S inducing a minimum-cost Nash equilibrium, which we denote by

$$\text{PoS}(G \mid S) = \max_{\substack{c \in \mathbb{R}_{\geq 0}^E : \\ S \text{ minimum-cost} \\ \text{Nash equilibrium}}} \frac{C(S)}{C(\sigma^*)}.$$

An additional simplification can be made while preserving a lower bound. The social optimum does not have to be exact. For any spanning tree $T \subseteq E$, we define

$$\text{PoS}(G \mid S, T) = \sup_{\substack{c \in \mathbb{R}_{\geq 0}^E : \\ S \text{ minimum-cost} \\ \text{Nash equilibrium}}} \frac{C(S)}{C(T)}.$$

Eventually, considering specific edge costs $c \in \mathbb{R}_{\geq 0}^E$ such that S induces a minimum-cost Nash equilibrium leaves a simple expression $C(S)/C(T)$ that we write as $\text{PoS}(G, c \mid S, T)$. The defined values constitute a chain of lower bounds on the price of stability PoS of the




Figure 5.15: The fan graph with n players.

problem, as we have

$$\text{PoS} \geq \text{PoS}(G) \geq \text{PoS}(G \mid S) \geq \text{PoS}(G \mid S, T) \geq \text{PoS}(G, c \mid S, T).$$

The choices of G , S , and T in the following are generally motivated by computational experiments. Choices resulting in good lower bounds seem to depend significantly on the cost function f . Therefore, we examine the price of stability for three different classes, which are constant total edge cost, affine total edge cost, and polynomial total edge cost. A common theme that appears is the structure S of a minimum-cost equilibrium. Computational experiments suggest that the best lower bounds are achieved by S being a star, that is, when every player connects to the root via a single edge. This makes sense on an intuitive level, a high price of stability originates from the lack of sharing edges in Nash equilibria.

5.4.3.1 The Fan Graph

For constant total edge cost and affine total edge cost, setting G to be the fan graph results in strong lower bounds according to computational experiments. The **fan graph** G_n on $n + 1$ vertices is depicted in Figure 5.15. Every vertex of a path on n vertices is connected by a direct edge to a common root r . Instances of this structure are first considered by Fiat et al. (2006) for the construction of lower bounds. Bilò, Caragiannis, et al. (2010, 2013) optimize the weights and get an improved lower bound.

We examine the price of stability $\text{PoS}(G_n \mid S_n)$ where S_n is the star in G_n , that is, the set of all edges that are incident to the root. The regularity of these instances leads to the first simplifications. The following lemmas show that we can focus on system optima that use exactly one edge that is incident to the root. These spanning trees are actually characterized by their single edge which is incident to the root. For $k \in [n]$, we denote by T_k the spanning tree that consists of $\{r, k\}$ and the path $1 - 2 - \dots - n$ in G_n .

Observation 5.35 (Decomposition with respect to the social optimum). *Let $T \subseteq E$ be a spanning tree in G_n such that $|\delta_T(r)| > 1$. Then*

$$\text{PoS}(G_n \mid S_n, T) \leq \max_{n' \in [n]} \max_{k \in [n']} \text{PoS}(G_{n'} \mid S_{n'}, T_k).$$

Proof. The tree T decomposes into $k \in \mathbb{N}$ trees $T^{(1)}, \dots, T^{(k)}$ such that $T = \bigcup_{i \in [k]} T^{(i)}$ and two trees $T^{(i)}$ and $T^{(j)}$, $i \neq j$ overlap in r only. Let $S^{(i)} := S_n[V(T^{(i)})]$ be the star induced by the vertex set of $T^{(i)}$ for every $i \in [k]$. Then $C(S_n)/C(T_n)$ is the median of $C(S^{(i)})/C(T^{(i)})$ over all $i \in [k]$. Now the statement follows from the median inequality. ■

A similar decomposition can be performed if there is an edge incident to the root which has cost zero. Due to the following lemma we can concentrate on instances with positive cost.

Observation 5.36 (Decomposition with respect to zero edges). *Let $c \in \mathbb{R}_{\geq 0}^{E_n}$ such that there is $i \in [n]$ with $c_{r,i} = 0$. Then there is a network cost-sharing game on a smaller fan graph with at least as good price of stability, that is,*

$$\text{PoS}(G_n, c \mid S_n) \leq \max_{k \in [n-1]} \text{PoS}(G_k \mid S_k).$$

Proof. Let $T \subseteq E$ be the support of a social optimum. Set \widehat{T} to be the result of T after contracting $\{r, i\}$ and replacing any occurrence of $\{i-1, i\}$ and $\{i, i+1\}$ by $\{r, i-1\}$ and $\{r, i+1\}$, respectively. By Lemma 5.4 and Observation 5.2, we may assume that \widehat{T} is again a tree. From S_n being a Nash equilibrium, we obtain $c_{r,i-1} < c_{i-1,i}$ and $c_{r,i+1} < c_{i,i+1}$. Therefore, $C(\widehat{T}) \leq C(T)$. Further, \widehat{T} does not contain $\{i-1, i\}$ or $\{i, i+1\}$. Now the statement follows from Observation 5.35. ■

The preceding two observations are based on decomposing instances on the fan graph into smaller ones. The following result shows the natural property that the price of stability on a fan graph is the higher the larger the graph is.

Lemma 5.37 (Monotonicity of the price of stability). *The price of stability in the fan is nondecreasing in its size, that is, for $n \leq n'$*

$$\text{PoS}(G_n \mid S_n) \leq \text{PoS}(G_{n'} \mid S_{n'}).$$

Proof. Let $c \in \mathbb{R}_{\geq 0}^{E_n}$ such that the star S_n is a cost-minimal Nash equilibrium in the broadcast game on G_n with cost c . Further, let $T \subseteq E_n$ be the support of a social optimum. We extend the game by one player. Set $c_{r,n+1} = 0$ and $c_{n-1,n} = \sum_{e \in E_n} c_e$. The exact value of $c_{n-1,n}$ is not important as long as it is large enough. Then S_{n+1} induces a Nash equilibrium and $T \cup \{\{r, n\}\}$ induces a social optimum in the broadcast game on G_{n+1} with the extended cost $c \in \mathbb{R}_{\geq 0}^{E_{n+1}}$. As the cost was chosen arbitrarily (under the condition that the star is a cost-minimal Nash equilibrium), we obtain $\text{PoS}(G_n \mid S_n) \leq \text{PoS}(G_{n+1} \mid S_{n+1})$. The statement follows by induction. ■

In order to obtain good bounds on $\text{PoS}(G_n \mid S_n)$, we need to understand the edge cost under which S_n indeed induces a minimum-cost Nash equilibrium. The following lemma gives a characterization for a class of cost functions f , which includes many natural cost functions (see Observation 5.39).

Lemma 5.38 (Stars as Nash equilibria). *Let f be a strictly decreasing player cost function such that*

$$\left(\frac{f(2)}{f(1)}\right)^{\frac{k}{2}} \leq \frac{f(k)}{f(1)} \text{ for all } k \in 2\mathbb{N}_{\geq 1} \quad \text{and} \quad \left(\frac{f(2)}{f(1)}\right)^{\frac{k-3}{2}} \frac{f(3)}{f(1)} \leq \frac{f(k)}{f(1)} \text{ for all } k \in 2\mathbb{N}_{\geq 1} + 1.$$

Further, let $c \in \mathbb{R}_{\geq 0}^{E_n}$ such that $c_{r,i} > 0$ for all $i \in [n]$. Then the star is a cost-minimal equilibrium with respect to c and f if and only if

$$(i) \quad c_{r,i} < c_{i-1,i} + \frac{f(2)}{f(1)}c_{r,i-1} \text{ for all } i = 2, \dots, n$$

$$(ii) \quad c_{r,i} < c_{i,i+1} + \frac{f(2)}{f(1)}c_{r,i+1} \text{ for all } i = 1, \dots, n-1$$

$$(iii) \quad \text{for all } i = 2, \dots, n-1$$

$$c_{r,i-1} < c_{i-1,i} + \frac{f(3)}{f(1)}c_{r,i} \quad \text{or}$$

$$c_{r,i+1} < c_{i,i+1} + \frac{f(3)}{f(1)}c_{r,i} \quad \text{or}$$

$$c_{r,i-1} + c_{r,i} + c_{r,i+1} \leq c_{i-1,i} + c_{i,i+1} + 3\frac{f(3)}{f(1)}c_{r,i}.$$

Proof. Sufficiency. Assume that (i), (ii), and (iii) hold. Let σ be the strategy profile induced by the star. It follows from (i) and (ii), that for all $i \in [n-1]$

$$c_\sigma(\sigma_{i+1} \setminus \sigma_i) - c'_\sigma(\sigma_i \setminus \sigma_{i+1}) = f(1)c_{r,i+1} - f(2)c_{r,i} < f(1)c_{i,i+1}$$

and

$$c_\sigma(\sigma_i \setminus \sigma_{i+1}) - c'_\sigma(\sigma_{i+1} \setminus \sigma_i) = f(1)c_{r,i} - f(2)c_{r,i+1} < f(1)c_{i,i+1}$$

hold. [Corollary 5.6](#) implies that σ is a Nash equilibrium.

Let $\bar{\sigma}$ be another equilibrium different from the star σ . Then there must be $j \in [n]$ such that $n_{\bar{\sigma}}(\{r, j\}) > 1$. Based on [Lemma 5.4](#), there are $1 \leq k \leq j \leq l \leq n$ such that the set of players using the edge $\{r, j\}$ is exactly $\{k, \dots, l\}$. Assume for contradiction that $j - k > l - j$, that is, k is further away from j than l . Then $l - k < 2(j - k)$. By combining (ii) for the indices $i = k, \dots, j-1$, we obtain

$$c_{r,k} < \sum_{i=1}^{j-k} \left(\frac{f(2)}{f(1)}\right)^{i-1} c_{k+i-1, k+i} + \left(\frac{f(2)}{f(1)}\right)^{j-k} c_{r,j}.$$

With the convention $f(0) = f(1)$, the monotonicity of f yields $f(2(i-1)) \leq f(i)$ for all $i \in \mathbb{N}_{\geq 1}$. When combining the assumption on f with this, we can continue the preceding

inequality to

$$\begin{aligned} f(1)c_{r,k} &< \sum_{i=1}^{j-k} f(2(i-1))c_{k+i-1,k+i} + f(2(j-k))c_{r,j} \\ &\leq \sum_{i=1}^{j-k} f(i)c_{k+i-1,k+i} + f(l-k+1)c_{r,j} = C_k(\bar{\sigma}). \end{aligned}$$

This means player k has an incentive to deviate to $\{r, k\}$, which contradicts $\bar{\sigma}$ being an equilibrium. Thus, the assumption $j - k > l - j$ is false. The case $j - k < l - j$ can be ruled out by symmetrical reasoning. Hence, $j - k = l - j$ has to hold.

We show that (iii) actually implies that always one of the first two conditions in (iii) has to hold. Assume that this is not the case; that is, there is $i \in \{2, \dots, n-1\}$ such that $f(1)c_{r,i-1} \geq f(1)c_{i-1,i} + f(3)c_{r,i}$ and $f(1)c_{r,i+1} \geq f(1)c_{i,i+1} + f(3)c_{r,i}$. As $c_{r,i} > 0$ and f decreases strictly, we get

$$c_{r,i-1} + c_{r,i} + c_{r,i+1} \geq c_{i-1,i} + c_{i,i+1} + \frac{2f(3) + f(1)}{f(1)}c_{r,i} > c_{i-1,i} + c_{i,i+1} + 3\frac{f(3)}{f(1)}c_{r,i}.$$

It follows that none of the inequalities in (iii) holds. This is a contradiction and we conclude that one of the first two inequalities in (iii) is valid.

Assume without loss of generality (due to symmetry) that for index j , the inequality $f(1)c_{r,j-1} < f(1)c_{j-1,j} + f(3)c_{r,j}$ holds. Similar to before, we combine this inequality with (i) for $i = k, \dots, j-2$ and obtain

$$c_{r,k} < \sum_{i=1}^{j-k} \left(\frac{f(2)}{f(1)}\right)^{i-1} c_{k+i-1,k+i} + \left(\frac{f(2)}{f(1)}\right)^{j-k-1} \frac{f(3)}{f(1)} c_{r,j}.$$

Again, we continue this inequality by using the assumptions on f and get

$$\begin{aligned} f(1)c_{r,k} &< \sum_{i=1}^{j-k} f(2(i-1))c_{k+i-1,k+i} + f(2(j-k)+1)c_{r,j} \\ &\leq \sum_{i=1}^{j-k} f(i)c_{k+i-1,k+i} + f(l-k+1)c_{r,j} = C_k(\bar{\sigma}). \end{aligned}$$

This shows that player k wants to deviate, which contradicts $\bar{\sigma}$ being a Nash equilibrium. In total, we get that under (i), (ii), and (iii), the star is the unique equilibrium. In particular, it minimizes the social cost among all Nash equilibria.

Necessity of (i) and (ii). For the converse direction, assume that the star σ is a cost-minimal equilibrium. Corollary 5.6 immediately implies that (i) and (ii) hold with weak inequality. Assume they do not hold in a strict sense. Then without loss of generality (due to symmetry), there is a player $i \in [n-1]$ such that $f(1)c_{r,i} = f(1)c_{i,i+1} + f(2)c_{r,i+1}$. Let

$\widehat{\sigma}_i = \{\{r, i+1\}, \{i+1, i\}\}$ be the strategy of player i via her neighbor $i+1$. Executing the improving dynamics on $(\sigma_{-i}, \widehat{\sigma}_i)$ results in a Nash equilibrium $\bar{\sigma}$ with smaller potential. We evaluate their respective potentials to obtain

$$C(\bar{\sigma}) \leq \Phi(\bar{\sigma}) \leq \Phi(\sigma_{-i}, \widehat{\sigma}_i) = \Phi(\sigma) - f(1)c_{r,i} + f(1)c_{i,i+1} + f(2)c_{r,i+1} = \Phi(\sigma) = C(\sigma).$$

As σ is a minimum-cost Nash equilibrium, equality has to hold. Hence, $\bar{\sigma} = (\sigma_{-i}, \widehat{\sigma}_i)$ follows. The change of the social cost under i deviating in σ to $\widehat{\sigma}_i$ is

$$\frac{C(\sigma_{-i}, \widehat{\sigma}_i) - C(\sigma)}{f(1)} = c_{i,i+1} - c_{r,i} + \left(\frac{2f(2)}{f(1)} - 1\right) c_{r,i+1} < c_{i,i+1} - c_{r,i} + \frac{f(2)}{f(1)} c_{r,i+1} = 0$$

as f is strictly decreasing. This, again, is a contradiction to the assumption that σ minimizes the social cost. It follows that (i) and (ii) hold (with strict inequality).

Necessity of (iii). Assume (iii) does not hold. Thus, there is a $j \in \{2, \dots, n-1\}$ such that all three constraints are violated. Define the strategy profile $\widehat{\sigma}$ which differs from σ in two strategies. Players $j-1$ and $j+1$ both play their respective paths via their common neighbor j ; that is, $\widehat{\sigma}_{j-1} = \{\{j-1, j\}, \{j, r\}\}$ and $\widehat{\sigma}_{j+1} = \{\{j+1, j\}, \{j, r\}\}$. We compare the social cost of $\widehat{\sigma}$ and σ to get

$$\frac{C(\widehat{\sigma}) - C(\sigma)}{f(1)} = c_{j-1,j} + c_{j,j+1} - c_{r,j-1} - c_{r,j+1} + \left(\frac{3f(3)}{f(1)} - 1\right) c_{r,j} < 0.$$

Therefore, $\widehat{\sigma}$ cannot be a Nash equilibrium.

As j violates (iii), we know in particular that $f(1)c_{r,j-1} \geq f(1)c_{j-1,j} + f(3)c_{r,j}$ and $f(1)c_{j,j+1} \geq f(1)c_{r,j+1} + f(3)c_{r,j}$. This means that players $j-1$ and $j+1$ do not want to deviate from $\widehat{\sigma}$ to their respective strategies in σ . It follows from [Corollary 5.6](#) that a player other than j has an incentive to deviate to a neighboring player in $[n] \setminus \{j-1, j, j+1\}$. Let $\bar{\sigma}$ be the resulting strategy profile after such a deviation. Note that it leads to $1 \leq k \leq i \leq l \leq n$ such that the set of players using the edge $\{r, i\}$ is exactly $\{k, \dots, l\}$ and $i-k \neq k-l$. In the proof of sufficiency of (i), (ii), and (iii), we show that under (i) and (ii), the player who deviated from $\widehat{\sigma}$ to σ prefers her direct edge to the root. This is a contradiction. \blacksquare

Observation 5.39 (Star equilibria under constant, affine, or polynomial total edge cost). *The requirements of [Lemma 5.38](#) on the cost function are met by*

(i) *constant total edge cost, that is, $f(k) = \frac{1}{k}$,*

(ii) *affine total edge cost, that is, $f(k) = \frac{1+(k-1)s}{k}$ for $s \in [0, 1]$, and*

(iii) *polynomial total edge cost, that is, $f(k) = k^{\alpha-1}$ for $\alpha \in [0, 1]$.*

Proof. Constant total edge cost (i) is a special case of affine total edge cost when choosing $s = 0$. In the following, we make use of the inequalities $2^k \geq 2k$ and $3 \cdot 2^{k-1} \geq 2k + 1$ for all $k \in \mathbb{N}$ which hold due to Bernoulli's inequality.

Affine total edge cost. For $k \in \mathbb{N}_{\geq 1}$, we use the binomial theorem and $s \leq 1$ to obtain

$$\left(\frac{f(2)}{f(1)}\right)^k = \left(\frac{1+s}{2}\right)^k = \frac{\sum_{i=0}^k \binom{k}{i} s^i}{2^k} \leq \frac{1+s \sum_{i=1}^k \binom{k}{i}}{2^k} = \frac{f(2^k)}{f(1)} \leq \frac{f(2k)}{f(1)}.$$

For $k \in \mathbb{N}_{\geq 1}$, we apply the same bound based on the binomial theorem twice to get

$$\frac{f(3)}{f(1)} \left(\frac{f(2)}{f(1)}\right)^{k-1} = \frac{s}{3} \left(\frac{f(2)}{f(1)}\right)^{k-1} + \frac{2}{3} \left(\frac{f(2)}{f(1)}\right)^k \leq \frac{s}{3} \frac{f(2^{k-1})}{f(1)} + \frac{2}{3} \frac{f(2^k)}{f(1)}.$$

Evaluating the right hand-side and using $s \leq 1$ gives

$$\frac{s}{3} \frac{f(2^{k-1})}{f(1)} + \frac{2}{3} \frac{f(2^k)}{f(1)} = \frac{1+s2^k + s^2(2^{k-1}-1)}{3 \cdot 2^{k-1}} \leq \frac{f(3 \cdot 2^{k-1})}{f(1)} \leq \frac{f(2k+1)}{f(1)}.$$

Polynomial total edge cost. For $k \in \mathbb{N}_{\geq 1}$, a straightforward calculation shows that

$$\left(\frac{f(2)}{f(1)}\right)^k = (2^k)^{\alpha-1} \leq (2k)^{\alpha-1} = \frac{f(2k)}{f(1)}$$

and

$$\frac{f(3)}{f(1)} \left(\frac{f(2)}{f(1)}\right)^{k-1} = (3 \cdot 2^k)^{\alpha-1} \leq (2k+1)^{\alpha-1} = \frac{f(2k+1)}{f(1)}. \quad \blacksquare$$

Observation 5.40. *From the first part of the proof of Lemma 5.38, we see that if the star is the best equilibrium, then the star is the unique equilibrium. Further, the last condition of Lemma 5.38 (iii) can be dropped.*

Due to Observation 5.35, we focus on determining $\text{PoS}(G_n | S_n, T_k)$. Based on Lemma 5.38, we examine the following program. Note that C depends implicitly on the decision variables c .

$$\begin{aligned} \max \quad & \frac{C(S_n)}{C(T_k)} && \text{(FAN)} \\ \text{s. t.} \quad & c_{r,i+1} < c_{i,i+1} + \frac{f(2)}{f(1)} c_{r,i} && \text{for all } i \in [n-1] \\ & c_{r,i} < c_{i,i+1} + \frac{f(2)}{f(1)} c_{r,i+1} && \text{for all } i \in [n-1] \\ & c_{r,i-1} < c_{i-1,i} + \frac{f(3)}{f(1)} c_{r,i} \text{ or } c_{r,i+1} < c_{i,i+1} + \frac{f(3)}{f(1)} c_{r,i} && \text{for all } i \in [n-1] \setminus \{1\} \\ & c_{i,i+1} \geq 0 && \text{for all } i \in [n-1]. \\ & c_{r,i} > 0 && \text{for all } i \in [n] \end{aligned}$$

We transform this program (FAN) to obtain a formulation that allows us to apply the theory of linear programming.

The set of feasible solutions to (FAN) is a cone. Scaling any feasible solution with a positive constant yields another feasible solution with the same objective function value. Hence, we can fix the denominator of the objective function to $f(1)$ and obtain an equivalent formulation.

Further, the strict inequalities in the program can be relaxed to weak inequalities. To see this, consider a feasible solution to this relaxation. Adding $\varepsilon > 0$ to all variables yields another feasible solution that fulfills all inequalities strictly. As the objective is continuous and ε can be chosen arbitrarily small, the relaxation is actually not proper; that is, changing the strict inequalities to weak inequalities preserves the optimal value.

Finally, we introduce binary variables $z \in \{0, 1\}^n$ representing the satisfied inequalities of the disjunctive constraints. Setting $z_i = 0$ corresponds to $c_{r,i-1} < c_{i-1,i} + \frac{f(3)}{f(1)}c_{r,i}$ and $z_i = 1$ to $c_{r,i+1} < c_{i,i+1} + \frac{f(3)}{f(1)}c_{r,i}$. Note that these inequalities dominate the corresponding inequalities in the first two sets of constraints in (FAN). Hence, we can combine them. To avoid case distinctions, we set $z_1 = 0$ and $z_n = 1$. For the sake of a more compact notation, we define the set of such vectors $Z_n := \{z \in \{0, 1\}^n \mid z_1 = 0, z_n = 1\}$.

Altogether, we obtain the equivalent program

$$\begin{aligned}
 \max \quad & \sum_{i=1}^n c_{r,i} && \text{(FAN')} \\
 \text{s. t.} \quad & 1 = \sum_{i=1}^{k-1} \frac{if(i)}{f(1)}c_{i,i+1} + \frac{nf(n)}{f(1)}c_{r,k} + \sum_{i=k}^{n-1} \frac{(n-i)f(n-i)}{f(1)}c_{i,i+1} \\
 & c_{r,i+1} \leq c_{i,i+1} + \frac{f(2+z_i)}{f(1)}c_{r,i} && \text{for all } i \in [n-1] \\
 & c_{r,i} \leq c_{i,i+1} + \frac{f(3-z_{i+1})}{f(1)}c_{r,i+1} && \text{for all } i \in [n-1] \\
 & c_{i,i+1} \geq 0 && \text{for all } i \in [n-1] \\
 & c_{r,i} \geq 0 && \text{for all } i \in [n] \\
 & z \in Z_n.
 \end{aligned}$$

Observation 5.41 (Symmetry of solutions). *If c, z is a solution to (FAN'), then also \bar{c}, \bar{z} is a solution where*

$$\begin{aligned}
 \bar{c}_{r,i} &= c_{r,n-i+1} && \text{for all } i \in [n] \\
 \bar{c}_{i,i+1} &= c_{n-i,n-i+1} && \text{for all } i \in [n-1] \\
 \bar{z}_i &= 1 - z_{n-i+1} && \text{for all } i \in [n]
 \end{aligned}$$

After fixing the variables z in (FAN'), we obtain a linear program. The following lemma regards one specific solution to it. The subsequent theorem shows in particular

that this solution determines $\text{PoS}(G_n \mid S_n, T_k)$. Together with the preceding findings, we obtain a characterization of $\text{PoS}(G_n \mid S_n)$.

Lemma 5.42 (Basic solutions). *Let f be a cost function fulfilling the requirements of Lemma 5.38. Fix some $z \in Z_n$. Then (FAN^z) has a unique positive basic feasible solution $c^{(z)}$, which is the solution to the system of equations*

$$\begin{aligned} c_{r,i+1} &= c_{i,i+1} + \frac{f(2+z_i)}{f(1)} c_{r,i} && \text{for all } i \in [n-1] \\ c_{r,i} &= c_{i,i+1} + \frac{f(3-z_{i+1})}{f(1)} c_{r,i+1} && \text{for all } i \in [n-1] \\ 1 &= \sum_{i=1}^{k-1} \frac{if(i)}{f(1)} c_{i,i+1} + \frac{nf(n)}{f(1)} c_{r,k} + \sum_{i=k}^{n-1} \frac{(n-i)f(n-i)}{f(1)} c_{i,i+1} \end{aligned}$$

Proof. We show that the system of equations defines a unique vector $c^{(z)}$ and check its feasibility and uniqueness afterwards.

A solution. Note that all but the last equation in the system are homogeneous; that is, they do not involve an additive constant. Hence, a solution satisfying all but the last equality can be scaled such that it fulfills the whole system. Consequently, we focus on the other equations. Eliminating the variables $c_{i,i+1}$ for all $i \in [n-1]$ yields

$$\frac{c_{r,i+1}}{c_{r,i}} = \frac{1 + \frac{f(2+z_i)}{f(1)}}{1 + \frac{f(3-z_{i+1})}{f(1)}} > 0 \quad \text{for all } i \in [n-1].$$

By substituting this into the original equations, we further obtain

$$\frac{c_{r,i}}{c_{i,i+1}} = \frac{1 + \frac{f(3-z_{i+1})}{f(1)}}{1 - \frac{f(3-z_{i+1})}{f(1)} \frac{f(2+z_i)}{f(1)}} > 0 \quad \text{for all } i \in [n-1].$$

These two families of equations clearly determine c up to scaling. Moreover, all components of c have the same sign (or are all zero). Since all coefficients of the last inequality in the defining system of c are positive, it follows that there is a unique solution $c^{(z)}$.

Feasibility and uniqueness. The system of $2n-1$ equations uniquely determines $2n-1$ variables. Hence, the equations must be linearly independent and $c^{(z)}$ is a basic solution of (FAN^z). As it is positive, it is also feasible. The remaining inequalities of (FAN^z) are nonnegativity constraints. Therefore, in every other basic solution there is a variable that is zero. ■

Theorem 5.43 (The price of stability in the fan). *Let f be a cost function fulfilling the requirements of Lemma 5.38. Further, let $c^{(z)}$ be the cost vector as defined in Lemma 5.42. Then the price of stability on the fan under the assumption that the star is a minimum-cost Nash equilibrium can be expressed as*

$$\text{PoS}(G_n \mid S_n) = \max_{n' \in [n]} \max_{k \in [n']} \max_{z \in Z_{n'}} \text{PoS}(G_{n'}, c^{(z)} \mid S_{n'}, T_k).$$

Proof. Due to the decomposition property shown in Observation 5.35, we obtain for the price of stability

$$\text{PoS}(G_n \mid S_n) = \max_{\substack{\text{spanning tree} \\ T \subseteq E_n}} \text{PoS}(G_n \mid S_n, T) = \max_{n' \in [n]} \max_{k \in [n']} \text{PoS}(G_{n'} \mid S_{n'}, T_k).$$

For $1 \leq k \leq n'$, the price of stability $\text{PoS}(G_{n'} \mid S_{n'}, T_k)$ can be determined using (FAN') due to Lemma 5.38. For a cost vector c with $c_{r,i} = 0$ for some $i \in [n']$, the characterization Lemma 5.38 does not hold. Still, Observation 5.36 shows that a smaller n' would achieve a larger or equal price of stability. After fixing the decisions $z \in \{0, 1\}^n$, the program (FAN') becomes linear. Thus, there is an optimal basic solution. Lemma 5.42 shows that $c^{(z)}$ is the unique positive basic feasible solution. Thus, any other basic feasible solution contains a variable that is zero. We show that all these other solutions can be reduced to smaller instances. If $c_{r,i} = 0$ for some $i \in [n]$, Observation 5.36 shows that the instance can be decomposed into smaller ones. If $c_{i,i+1} = 0$ for some $i \in [n-1]$, then the constraints $c_{r,i+1} \leq c_{i,i+1} + f(2+z_i)/f(1) \cdot c_{r,i}$ and $c_{r,i} \leq c_{i,i+1} + f(3-z_{i+1})/f(1) \cdot c_{r,i+1}$ yield $c_{r,i} = c_{r,i+1} = 0$. Thus, also in this case the instance decomposes. We therefore can focus on the solutions $c^{(z)}$ and get

$$\begin{aligned} \text{PoS}(G_n \mid S_n) &= \max_{n' \in [n]} \max_{k \in [n']} \max_{z \in Z_{n'}} \text{val}(\text{FAN}') \\ &= \max_{n' \in [n]} \max_{k \in [n']} \max_{z \in Z_{n'}} \text{PoS}(G_{n'}, c^{(z)} \mid S_{n'}, T_k). \quad \blacksquare \end{aligned}$$

The objective function value of the feasible basic solution corresponding to $c^{(z)}$, z for some $z \in Z_n$ evaluates to

$$\frac{\sum_{i=1}^n c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} \frac{if(i)}{f(1)} c_{i,i+1}^{(z)} + \frac{nf(n)}{f(1)} c_{r,k}^{(z)} + \sum_{i=k}^{n-1} \frac{(n-i)f(n-i)}{f(1)} c_{i,i+1}^{(z)}}.$$

This fraction can be interpreted as the mediant of the fractions

$$\frac{c_{r,1}^{(z)}}{c_{1,2}^{(z)}}, \dots, \frac{c_{r,k-1}^{(z)}}{\frac{(k-1)f(k-1)}{f(1)} c_{k-1,k}^{(z)}}, \frac{c_{r,k}^{(z)}}{\frac{nf(n)}{f(1)} c_{r,k}^{(z)}}, \frac{c_{r,k+1}^{(z)}}{\frac{(n-k)f(n-k)}{f(1)} c_{k,k+1}^{(z)}}, \dots, \frac{c_{r,n}^{(z)}}{c_{n-1,n}^{(z)}}.$$

The first $k-1$ fractions and the last $n-k$ fractions are associated with the two branches of the tree T_k , respectively. The k -th term corresponds to its stem. Due to the mediant

inequality, the price of stability cannot be greater than the largest of these fractions. Hence, the following ratios play an important role for determining (a lower bound on) the price of stability. As deduced in the proof of [Lemma 5.42](#), the inequalities defining $c^{(z)}$ yield for $i \in [n-1]$

$$\frac{c_{r,i+1}^{(z)}}{c_{r,i}^{(z)}} = \frac{1 + \frac{f(2+z_i)}{f(1)}}{1 + \frac{f(3-z_{i+1})}{f(1)}}, \quad \frac{c_{r,i}^{(z)}}{c_{i,i+1}^{(z)}} = \frac{1 + \frac{f(3-z_{i+1})}{f(1)}}{1 - \frac{f(3-z_{i+1})}{f(1)} \frac{f(2+z_i)}{f(1)}}, \quad \frac{c_{r,i+1}^{(z)}}{c_{i,i+1}^{(z)}} = \frac{1 + \frac{f(2+z_i)}{f(1)}}{1 - \frac{f(2+z_i)}{f(1)} \frac{f(3-z_{i+1})}{f(1)}}.$$

The possible values that the latter two ratios can assume are

$$\frac{3}{2} \leq \frac{1}{1 - \frac{f(3)}{f(1)}} = \frac{1 + \frac{f(3)}{f(1)}}{1 - \left(\frac{f(3)}{f(1)}\right)^2} \leq \frac{1 + \frac{f(3)}{f(1)}}{1 - \frac{f(3)}{f(1)} \frac{f(2)}{f(1)}} \leq \frac{1 + \frac{f(2)}{f(1)}}{1 - \frac{f(2)}{f(1)} \frac{f(3)}{f(1)}} \leq \frac{1 + \frac{f(2)}{f(1)}}{1 - \left(\frac{f(2)}{f(1)}\right)^2} = \frac{1}{1 - \frac{f(2)}{f(1)}}.$$

This order is due to f being nonincreasing. As further $k \mapsto kf(k)$ is nondecreasing, $1/(1 - f(2)/f(1))$ is an upper bound to the objective function value obtained from $c^{(z)}, z$. We obtain the following corollary to [Theorem 5.43](#).

Corollary 5.44. *Let f be a cost function fulfilling the requirements of [Lemma 5.38](#). Then for every $n \in \mathbb{N}$, the price of stability on the fan under the assumption that the star is a cost-minimal Nash equilibrium is at most*

$$\text{PoS}(G_n \mid S_n) \leq \frac{f(1)}{f(1) - f(2)}.$$

The smallest ratio contributing to the objective function interpreted as mediant is $(f(1)c_{r,k}^{(z)})/(nf(n)c_{r,k}^{(z)})$. Therefore, it is beneficial for obtaining a large price of stability to put more weight on the other ratios. On the one hand, this can be achieved through the choice of z . As we see in the next sections, thereby we can get an exponential growth of the weight of the i -th term in the distance $|i - k|$ to the stem. On the other hand, increasing n includes more terms in the mediant with ratio at least $\frac{3}{2}$. It has, however, a converse effect. As $k \mapsto kf(k)$ is nondecreasing, the larger n is the smaller is the ratio of the terms with index close to k . It depends on the cost function f which of these two opposite influences is stronger.

5.4.3.2 A Lower Bound for Constant Total Edge Cost

For constant total edge cost, we obtain $\text{PoS}(G_n \mid S_n) \leq 2$ from [Corollary 5.44](#). We continue with a more careful analysis in order to determine $\text{PoS}(G_n \mid S_n)$ exactly for $f(k) = 1/k$. According to [Theorem 5.43](#), it suffices to understand which $z \in Z_n$ maximizes the value of (FAN'). To that end, we interpret the objective function value attained by $c^{(z)}$ and z as the harmonic mean of the ratios

$$\frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)}}, \quad \frac{c_{r,k}^{(z)}}{c_{r,k}^{(z)}}, \quad \frac{\sum_{i=k+1}^n c_{r,i}^{(z)}}{\sum_{i=k+1}^n c_{i-1,i}^{(z)}}$$

$\frac{c_{r,i+1}^{(z)}}{c_{r,i}^{(z)}}$	$z_{i+1} = 0$	$z_{i+1} = 1$	$\frac{c_{r,i+1}^{(z)}}{c_{i,i+1}^{(z)}}$	$z_{i+1} = 0$	$z_{i+1} = 1$
$z_i = 0$	$\frac{9}{8}$	1	$z_i = 0$	$\frac{9}{5}$	2
$z_i = 1$	1	$\frac{8}{9}$	$z_i = 1$	$\frac{3}{2}$	$\frac{8}{5}$

Table 5.1: The solution $c^{(z)}$ under the cost function $f(k) = 1/k$.

weighted by $\sum_{i=1}^{k-1} c_{r,i}^{(z)}$, $c_{r,k}^{(z)}$, and $\sum_{i=k+1}^n c_{r,i}^{(z)}$, respectively. We focus on the third ratio and its weight. By symmetry considerations along the lines of [Observation 5.41](#), the results extend to the first ratio and its weight. [Table 5.1](#) contains the relevant ratios of the costs of pairs of single edges for constant total edge cost that we use in the following.

Lemma 5.45 (Optimal choice of z with $z_1 = 0$). *Among all $z \in Z_n$, the choice $z = \mathbb{1}_n$ simultaneously maximizes the two ratios*

$$\frac{\sum_{i=2}^n c_{r,i}^{(z)}}{c_{r,1}^{(z)}} \quad \text{and} \quad \frac{\sum_{i=2}^n c_{r,i}^{(z)}}{\sum_{i=2}^n c_{i-1,i}^{(z)}}.$$

Proof. Consulting [Table 5.1](#) shows that the former ratio is maximized for $z = \mathbb{1}_n$. Hence, we focus on the latter ratio in the following. Assume for a contradiction that $z = \mathbb{1}_n$ does not attain the maximum. Fix $z \in Z_n$ to be a maximizer of the ratio at hand. Let $j \in \mathbb{N}$ be the smallest index such that $z_j = 1$. By our assumption, we have $j < n$. Let $\bar{z} = z - \mathbb{1}_j$ be equal to z except for $\bar{z}_j = 0$. We show that \bar{z} has a better ratio than z which contradicts its choice. For that purpose, we interpret the ratio of z as the harmonic mean of the two ratios

$$\frac{\sum_{i=2}^{j+1} c_{r,i}^{(z)}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(z)}} \quad \text{and} \quad \frac{\sum_{i=j+2}^n c_{r,i}^{(z)}}{\sum_{i=j+2}^n c_{i-1,i}^{(z)}}$$

weighted by $\sum_{i=2}^{j+1} c_{r,i}^{(z)}$ and $\sum_{i=j+2}^n c_{r,i}^{(z)}$, respectively.

Comparing the two ratios. First, we show that the first ratio is less than the second ratio. We evaluate the ratio of the terms up to the index $j+1$ using the values from [Table 5.1](#). In the case $z_{j+1} = 0$, we get

$$\frac{\sum_{i=2}^{j+1} c_{r,i}^{(z)}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(z)}} = \frac{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{j-2} + \left(\frac{9}{8}\right)^{j-2}}{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{j-2} \frac{1}{2} + \left(\frac{9}{8}\right)^{j-2} \frac{2}{3}} = \frac{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{j-2} 2}{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{j-2} \frac{7}{6}}.$$

As $\frac{12}{7} < \frac{9}{5}$, the properties of the harmonic mean imply that the ratio of the terms up to the index $j+1$ is strictly less than $\frac{9}{5}$. If $z_{j+1} = 1$, this ratio is

$$\frac{\sum_{i=2}^{j+1} c_{r,i}^{(z)}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(z)}} = \frac{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{j-2} + \left(\frac{9}{8}\right)^{j-3}}{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{j-2} \frac{1}{2} + \left(\frac{9}{8}\right)^{j-3} \frac{5}{8}} = \frac{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{j-2} \frac{17}{9}}{\sum_{i=2}^{j-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{j-2} \frac{19}{18}}.$$

Again this ratio is below $\frac{9}{5}$ because $\frac{34}{19} < \frac{9}{5}$.

For the choice $\mathbb{1}_n$, however, we obtain a ratio of

$$\frac{\sum_{i=2}^n c_{r,i}^{(\mathbb{1}_n)}}{\sum_{i=2}^n c_{i-1,i}^{(\mathbb{1}_n)}} = \frac{\sum_{i=2}^{n-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{n-2}}{\sum_{i=2}^{n-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{n-2} \frac{1}{2}} = \frac{10 \left(\frac{9}{8}\right)^{n-2} - 9}{\frac{11}{2} \left(\frac{9}{8}\right)^{n-2} - 5} \searrow \frac{20}{11} \text{ for } n \rightarrow \infty.$$

As z was chosen to be optimal, its overall ratio has to be larger than $\frac{20}{11}$. The mediant inequality yields

$$\frac{\sum_{i=2}^{j+1} c_{r,i}^{(z)}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(z)}} < \frac{9}{5} < \frac{20}{11} < \frac{\sum_{i=j+2}^n c_{r,i}^{(z)}}{\sum_{i=j+2}^n c_{i-1,i}^{(z)}}.$$

Comparing to \bar{z} . Recomputing the ratios for \bar{z} shows

$$\frac{\sum_{i=2}^{j+1} c_{r,i}^{(z)}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(z)}} < \frac{9}{5} = \frac{\sum_{i=2}^{j+1} c_{r,i}^{(\bar{z})}}{\sum_{i=2}^{j+1} c_{i-1,i}^{(\bar{z})}} \quad \text{and} \quad \frac{\sum_{i=j+2}^n c_{r,i}^{(z)}}{\sum_{i=j+2}^n c_{i-1,i}^{(z)}} = \frac{\sum_{i=j+2}^n c_{r,i}^{(\bar{z})}}{\sum_{i=j+2}^n c_{i-1,i}^{(\bar{z})}}.$$

Additionally, we get that the weight of the second part, which has the higher ratio, increases by setting $\bar{z}_j = 0$ relative to the weight of the first part, as

$$\sum_{i=2}^{j+1} c_{r,i}^{(\bar{z})} < \left(\frac{9}{8}\right)^2 \sum_{i=2}^{j+1} c_{r,i}^{(z)} \quad \text{and} \quad \sum_{i=j+2}^n c_{r,i}^{(\bar{z})} = \left(\frac{9}{8}\right)^2 \sum_{i=j+2}^n c_{r,i}^{(z)}.$$

In total, the monotonicity of the weighted harmonic mean implies that \bar{z} has a larger total ratio than z . This is the contradiction we aimed for. \blacksquare

Lemma 5.46 (Optimal choice of z with $z_1 = 1$). *Among all $z \in \{0, 1\}^n$ with $z_1 = 1$ and $z_n = 1$, the choice $z = \mathbb{1}_1 + \mathbb{1}_n$ simultaneously maximizes the two ratios*

$$\frac{\sum_{i=2}^n c_{r,i}^{(z)}}{c_{r,1}^{(z)}} \quad \text{and} \quad \frac{\sum_{i=2}^n c_{r,i}^{(z)}}{\sum_{i=2}^n c_{i-1,i}^{(z)}}.$$

Proof. Again, [Table 5.1](#) shows that the first ratio is maximized by $z = \mathbb{1}_1 + \mathbb{1}_n$. Thus, we focus on the second ratio. Let $z \in \{0, 1\}^n$ such that $z_1 = 1$ and $z_n = 1$. Further, let $1 = i_0 < i_1 < \dots < i_k = n$ be the indices for which z is one; that is, $z_i = 1$ if and only if

$i \in \{i_0, \dots, i_k\}$. We interpret the ratio in question for z as the mediant of the k ratios

$$\frac{\sum_{i=i_{j-1}+1}^{i_j} c_{r,i}^{(z)}}{\sum_{i=i_{j-1}+1}^{i_j} c_{i-1,i}^{(z)}}, j \in [k].$$

Let $j \in [k]$ and $l = i_j - i_{j-1}$ be the number of summands in the respective ratio. Its evaluation yields

$$\frac{\sum_{i=i_{j-1}+1}^{i_j} c_{r,i}^{(z)}}{\sum_{i=i_{j-1}+1}^{i_j} c_{i-1,i}^{(z)}} = \frac{1 + \sum_{i=2}^{l-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{l-2}}{\frac{2}{3} + \sum_{i=2}^{l-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{l-2} \frac{1}{2}} = \frac{10 \left(\frac{9}{8}\right)^{l-2} - 8}{\frac{11}{2} \left(\frac{9}{8}\right)^{l-2} - \frac{13}{3}} \nearrow \frac{20}{11} \text{ for } l \rightarrow \infty.$$

Note that even though the first equation is based on $l > 1$, the overall evaluation is the same for $l = 1$. As this is increasing in l , the mediant inequality shows that $k = 1$ and, therefore, $z = \mathbb{1}_1 + \mathbb{1}_n$ maximizes the total ratio. \blacksquare

Lemma 5.47 (Price of stability in the fan under constant total edge cost). *Under the assumption that the star is a minimum-cost Nash equilibrium, the price of stability in the fan is $\frac{20}{11}$. More specifically, for fixed number of players n , the worst-case is attained by*

$$\text{PoS}(G_n \mid S_n) = \text{PoS}(G_n, c^{(1_n)} \mid S_n, T_1) = \frac{20 \left(\frac{9}{8}\right)^{n-2} - 16}{11 \left(\frac{9}{8}\right)^{n-2} - 8} \nearrow \frac{20}{11} \text{ for } n \rightarrow \infty.$$

Proof. [Theorem 5.43](#) guarantees the existence of $n' \in [n]$, $k \in [n']$, $z \in Z_{n'}$ such that

$$\text{PoS}(G_n \mid S_n) = \text{PoS}(G_{n'}, c^{(z)} \mid S_{n'}, T_k) = \frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)} + c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)} + c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}}.$$

We first argue that $z_i = 1$ for all $1 < i < k$ and $z_i = 0$ for all $k < i < n$ based on the monotonicity of certain harmonic means. Then we finish the proof by comparing the explicit evaluations for the possible values of n' and k .

We interpret the price of stability attained by n' , k , and z as the harmonic mean of the three ratios

$$\frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)}}, \quad \frac{c_{r,k}^{(z)}}{c_{r,k}^{(z)}}, \quad \text{and} \quad \frac{\sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{\sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}}$$

weighted by $\sum_{i=1}^{k-1} c_{r,i}^{(z)}$, $c_{r,k}^{(z)}$, and $\sum_{i=k+1}^{n'} c_{r,i}^{(z)}$, respectively. Note that these ratios correspond to the stem and the two branches of T_k . If a branch does not exist due to $k = 1$ or $k = n'$, we treat the associated ratio as one.

The high-ratio branch. Due to [Observation 5.41](#), we can assume without loss of generality that these three ratios are ordered by

$$1 = \frac{c_{r,k}^{(z)}}{c_{r,k}^{(z)}} \leq \frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)}} \leq \frac{\sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{\sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}}.$$

From [Lemmas 5.45](#) and [5.46](#), we know that $z_i = 0$ for all $k < i < n'$ simultaneously maximizes the ratio on the right hand-side and its weight relative to the other weights. From the monotonicity of the weighted harmonic mean, it follows that this also maximizes the price of stability attained by n' , k , and z . Thus, we may assume $z_i = 0$ for all $k < i < n'$.

The low-ratio branch. We slightly reinterpret the price of stability as the harmonic mean of the two ratios given by the low-ratio branch and the combination of the stem with the high-ratio branch. Assume for contradiction that the inequality

$$\frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)}} \geq \frac{c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}},$$

does not hold. Then, the mediant inequality implies

$$\frac{\sum_{i=1}^{k-1} c_{r,i}^{(z)} + c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{\sum_{i=1}^{k-1} c_{i,i+1}^{(z)} + c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}} < \frac{c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{r,i}^{(z)}}{c_{r,k}^{(z)} + \sum_{i=k+1}^{n'} c_{i-1,i}^{(z)}}.$$

Note that the right-hand side almost corresponds to the price of stability of the instance that is obtained when deleting the vertices $1, \dots, k-1$, which is another fan. The only difference might be that $z_k = 1$. The computations in the proofs of [Lemmas 5.45](#) and [5.46](#) reveal that $z_k = 0$ and, therefore, the smaller instance would give an even larger ratio. This contradicts the optimal choice of n' , k , and z . Hence, the assumption is wrong and the inequality holds. Then, repeating the reasoning from the high-ratio branch under the additional use of [Observation 5.41](#) shows that we may assume $z_i = 1$ for all $1 < i < k$.

The stem. Now, we can compute $\text{PoS}(G_{n'}, c^{(z)} \mid S_{n'}, T_k)$ under the assumption that $z_i = 1$ for $1 < i < k$ and $z_i = 0$ for $k < i < n$. We invoke [Observation 5.41](#) once again to see that we may assume $z_k = 0$ as well as $k < n$. The costs are depicted in [Figure 5.16](#). For $k > 2$, $\text{PoS}(G_{n'}, c^{(z)} \mid S_{n'}, T_k)$ evaluates to

$$\frac{\left(\frac{9}{8}\right)^{k-3} + \sum_{i=2}^{k-2} \left(\frac{9}{8}\right)^{k-i-1} + 1 + 1 + \sum_{i=k+1}^{n-1} \left(\frac{9}{8}\right)^{i-k} + \left(\frac{9}{8}\right)^{n-k-1}}{\left(\frac{9}{8}\right)^{k-3} \frac{1}{2} + \sum_{i=2}^{k-2} \left(\frac{9}{8}\right)^{k-i-1} \frac{5}{9} + \frac{2}{3} + 1 + \sum_{i=k+1}^{n-1} \left(\frac{9}{8}\right)^{i-k} \frac{5}{9} + \left(\frac{9}{8}\right)^{n-k-1} \frac{1}{2}}.$$

For $k = 2$, we obtain

$$\frac{\left(\frac{9}{8}\right)^{-1} + 1 + \sum_{i=3}^{n-1} \left(\frac{9}{8}\right)^{i-2} + \left(\frac{9}{8}\right)^{n-3}}{\left(\frac{9}{8}\right)^{-1} \frac{5}{8} + 1 + \sum_{i=3}^{n-1} \left(\frac{9}{8}\right)^{i-2} \frac{5}{9} + \left(\frac{9}{8}\right)^{n-3} \frac{1}{2}}.$$

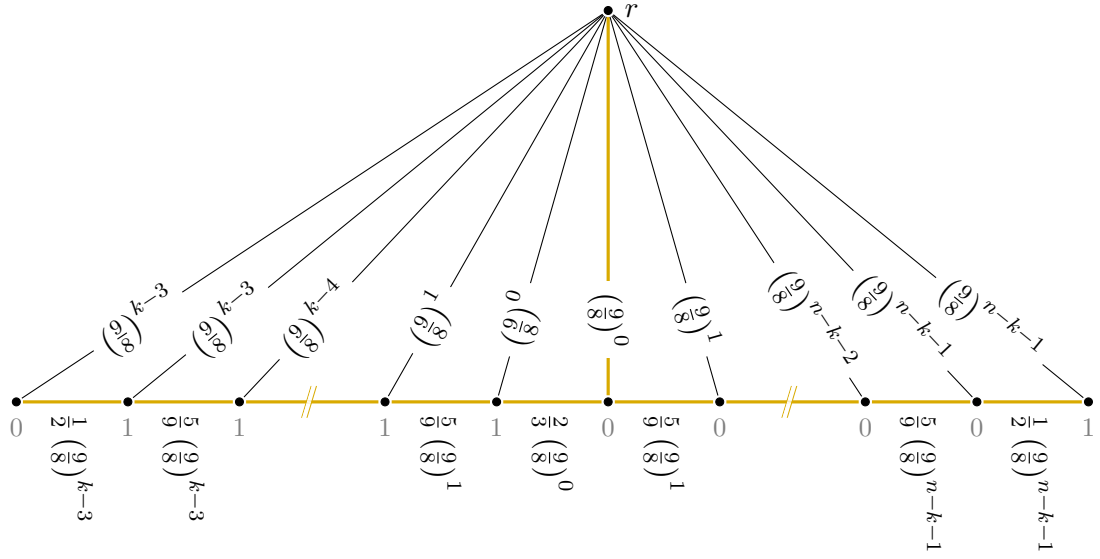


Figure 5.16: The edge cost $c^{(z)}$ in the fan with $z_i = 1$ for $1 < i < k$ and $z_i = 0$ for $k \leq i < n$. The vertices $i \in [n]$ are annotated with their value of z_i in gray. T_k is highlighted in yellow.

In both these cases, the fractions are equal to

$$\frac{10 \left(\left(\frac{9}{8}\right)^{k-3} + \left(\frac{9}{8}\right)^{n-k-1} \right) - 16}{\frac{11}{2} \left(\left(\frac{9}{8}\right)^{k-3} + \left(\frac{9}{8}\right)^{n-k-1} \right) - \frac{25}{3}} \nearrow \frac{20}{11} \text{ for } n \rightarrow \infty.$$

For $k = 1$, on the other hand we obtain

$$\frac{1 + \sum_{i=2}^{n-1} \left(\frac{9}{8}\right)^{i-1} + \left(\frac{9}{8}\right)^{n-2}}{1 + \sum_{i=2}^{n-1} \left(\frac{9}{8}\right)^{i-1} \frac{5}{9} + \left(\frac{9}{8}\right)^{n-2} \frac{1}{2}} = \frac{10 \left(\frac{9}{8}\right)^{n-2} - 8}{\frac{11}{2} \left(\frac{9}{8}\right)^{n-2} - 4} \nearrow \frac{20}{11} \text{ for } n \rightarrow \infty.$$

It can be checked that for fixed n' the value is maximized by $k = 1$. As for $k = 1$ the value is increasing in n' , it follows $n' = n$. \blacksquare

5.4.3.3 A Lower Bound for Affine Total Edge Cost

Computational experiments suggest that fan graphs result in best possible lower bounds for affine total edge cost as well. To obtain suitable cost vectors $c^{(z)}$, we use the choice for $z \in Z_n$ that proved optimal for constant total edge cost. For $n \in \mathbb{N}$ and $k \in [n]$, let $z^k \in Z_n$ be the vector such that $z_i = 1$ if and only if $1 < i < k$ or $i = n$. We evaluate the lower bound

$$\text{PoS}(G_n \mid S_n) \geq \max_{n' \in [n]} \max_{k \in [n']} \text{PoS}(G_{n'}, c^{(z^k)} \mid S_{n'}, T_k)$$

$\frac{c_{r,i+1}^{(z)}}{c_{r,i}^{(z)}}$	$z_{i+1} = 0$	$z_{i+1} = 1$	$\frac{c_{r,i+1}^{(z)}}{c_{i,i+1}^{(z)}}$	$z_{i+1} = 0$	$z_{i+1} = 1$
$z_i = 0$	$\frac{3(3+s)}{4(2+s)}$	1	$z_i = 0$	$\frac{3(3+s)}{(1-s)(2s+5)}$	$\frac{2}{1-s}$
$z_i = 1$	1	$\frac{4(2+s)}{3(3+s)}$	$z_i = 1$	$\frac{3}{2(1-s)}$	$\frac{4(2+s)}{(1-s)(2s+5)}$

Table 5.2: The solution $c^{(z)}$ under the cost function $f(k) = s + (1-s)/k$.

that is implied by [Theorem 5.43](#). We use the ratios in [Table 5.2](#) to evaluate the social cost of the star under this choice of z . For $2 < k < n$, we obtain

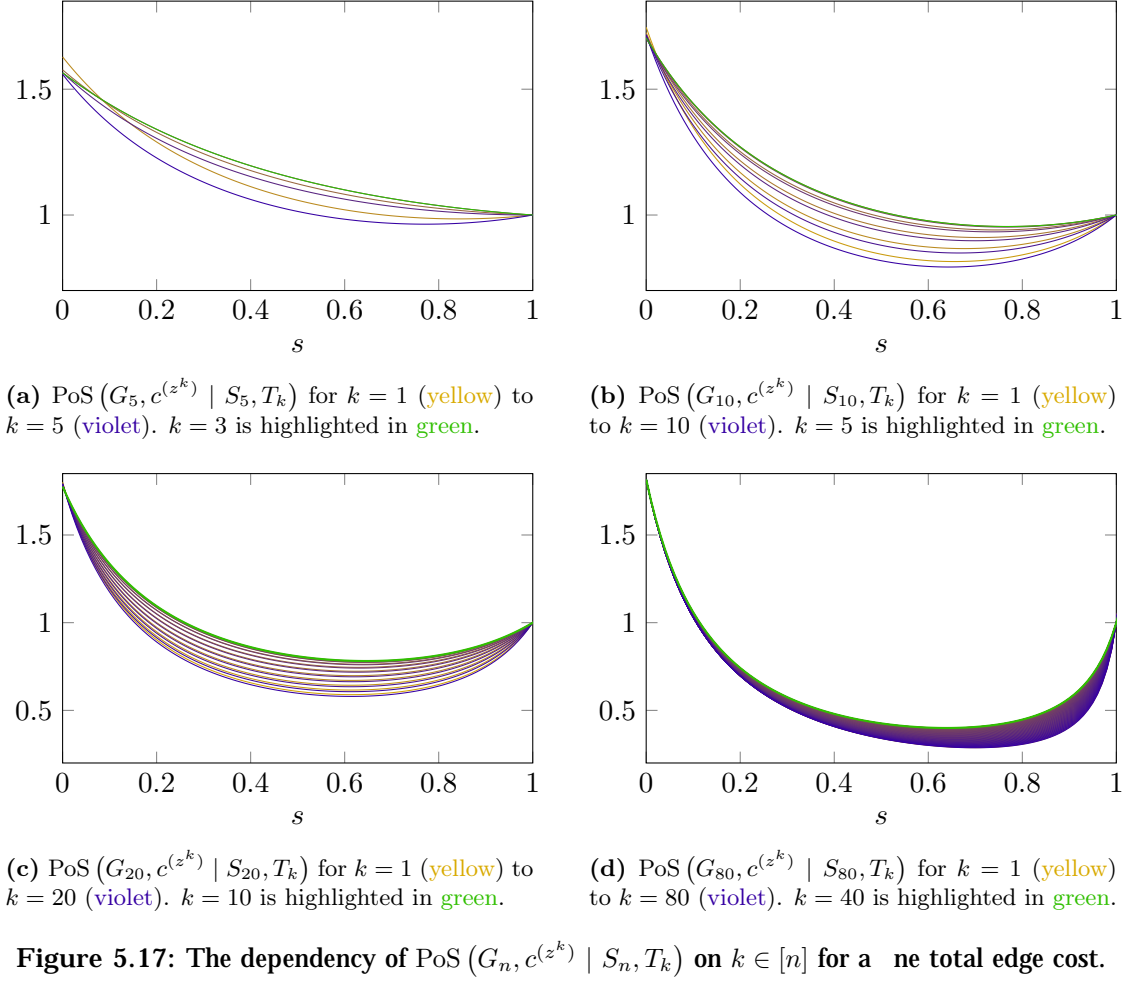
$$\frac{C(S_n)}{c_{r,k}^{(z^k)}} = \left(\frac{3(3+s)}{4(2+s)}\right)^{k-3} + \sum_{i=2}^{k-1} \left(\frac{3(3+s)}{4(2+s)}\right)^{k-1-i} + 1 + \sum_{i=k+1}^{n-1} \left(\frac{3(3+s)}{4(2+s)}\right)^{i-k} + \left(\frac{3(3+s)}{4(2+s)}\right)^{n-k-1}$$

Simplifying the right-hand side and doing the computations for the extreme cases of k shows for $n > 1$

$$\frac{C(S_n)}{c_{r,k}^{(z^k)}} = \begin{cases} \frac{2(5+s)}{1-s} \left(\frac{3(3+s)}{4(2+s)}\right)^{n-k-1} - \frac{8+4s}{1-s} & \text{if } k = 1 \\ \frac{2(5+s)}{1-s} \left(\left(\frac{3(3+s)}{4(2+s)}\right)^{k-3} + \left(\frac{3(3+s)}{4(2+s)}\right)^{n-k-1} \right) - \frac{8(2+s)}{1-s} & \text{if } 1 < k < n \\ \frac{2(5+s)}{1-s} \left(\frac{3(3+s)}{4(2+s)}\right)^{k-3} - \frac{7+5s}{1-s} & \text{if } k = n. \end{cases}$$

The cost of a social optimum on the other hand for $2 < k < n$ is

$$\begin{aligned} \frac{C(T_k)}{c_{r,k}^{(z^k)}} &= \left(\frac{3(3+s)}{4(2+s)}\right)^{k-3} \frac{2}{1-s} + \sum_{i=2}^{k-2} (1+s(i-1)) \left(\frac{3(3+s)}{4(2+s)}\right)^{k-1-i} \frac{3(3+s)}{(1-s)(2s+5)} \\ &+ (1+s(k-2)) \frac{3}{2(1-s)} + (1+s(n-1)) \\ &+ \sum_{i=k+1}^{n-1} (1+s(n-i)) \left(\frac{3(3+s)}{4(2+s)}\right)^{i-k} \frac{3(3+s)}{(1-s)(2s+5)} + \left(\frac{3(3+s)}{4(2+s)}\right)^{n-k-1} \frac{2}{1-s} \end{aligned}$$



Doing the calculations also for the extreme cases of k yields for $n > 1$

$$\frac{C(T_k)}{c_{r,k}^{(z)}} = \begin{cases} \left(\frac{3(3+s)}{4(2+s)} \right)^{n-k-1} \frac{12s^3+63s^2+82s+11}{2(1-s)} - 2s(s+2)n + s(2s+5)k - \frac{8s^3+33s^2+39s+4}{1-s} & \text{if } k = 1 \\ \left(\left(\frac{3(3+s)}{4(2+s)} \right)^{k-3} + \left(\frac{3(3+s)}{4(2+s)} \right)^{n-k-1} \right) \frac{12s^3+63s^2+82s+11}{2(1-s)} - 2s(s+2)n + \frac{2s(1-s)}{3}k - \frac{64s^3+209s^2+206s+25}{3(1-s)} & \text{if } 1 < k < n \\ \left(\frac{3(3+s)}{4(2+s)} \right)^{k-3} \frac{12s^3+63s^2+82s+11}{2(1-s)} + sn - \frac{s(8s+13)}{3}k - \frac{40s^3+107s^2+95s+10}{3(1-s)} & \text{if } k = n. \end{cases}$$

For fixed $n \in \mathbb{N}$, the choice of the social optimum that maximizes the bound on the price of stability depends on s . Figure 5.17 shows the dependency of PoS ($G_n, c^{(z^k)} \mid S_n, T_k$)




Figure 5.18: Lower bounds to the price of stability for a ne total edge cost. The bounds $\text{PoS}(G_n, c^{(z_{\lceil n/2 \rceil})} \mid S_n, T_{\lceil n/2 \rceil})$ are plotted for $n = 1$ (yellow) to $n = 100$ (violet). The bound resulting from maximizing over all $n \in \mathbb{N}$ is highlighted in green. The upper bound based on Theorem 5.23 is plotted in red.

on k for $n \in \{5, 10, 20, 80\}$. The diagrams show that our bounds can be below one. This results from the fact that T_k is not enforced to be a social optimum in (FAN'). The choice $k = \lceil n/2 \rceil$ generally results in good lower bounds on the price of stability for fixed n . Only for small $s \geq 0$, setting $k = 1$ yields better results. Note that this is in accordance with Lemma 5.47, which applies to $s = 0$. This latter effect becomes irrelevant when maximizing over all $n \in \mathbb{N}$. Then the improved bound from $k = 1$ is dominated by increasing n and using $k = \lceil n/2 \rceil$ again. Therefore, we focus on the lower bound

$$\text{PoS}(G_n \mid S_n) \geq \max_{n' \in [n]} \text{PoS}(G_{n'}, c^{(z_{\lceil n/2 \rceil})} \mid S_{n'}, T_{\lceil n/2 \rceil}).$$

It is visualized in Figure 5.18. For $s = 0$, the limit for $n \rightarrow +\infty$ is the same when choosing $k = 1$ or $k = \lceil n/2 \rceil$ (or any other k). Hence, we recover $\text{PoS} \geq 20/11$ in that case as already shown by Lemma 5.47. Interestingly, the best lower bound for $s > 0$ is obtained by a finite $n \in \mathbb{N}$. The reason becomes evident when consulting the interpretation of the price of stability as the mediant of ratios from pairs of edges, (see the end of Section 5.4.3.1). The increase of n includes more ratios in the mediant that

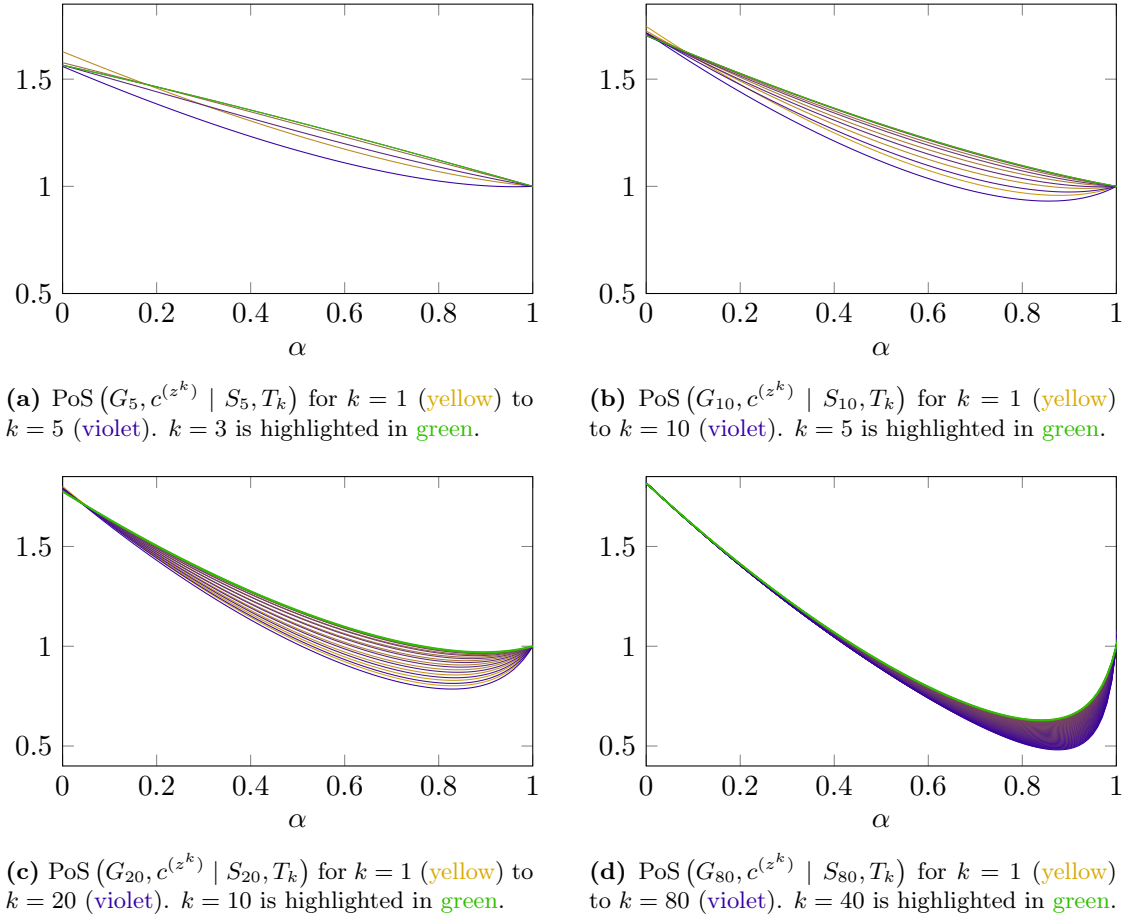


Figure 5.19: The dependency of PoS ($G_n, c^{(z^k)} \mid S_n, T_k$) on $k \in [n]$ for polynomial total edge cost.

make up for the bad ratio of the stem of T_k . The ratio of these additional pairs, however, gets smaller with n . At the same time, the ratio of the stem decreases. For large s , this gain from increasing n quickly turns into a loss. Thus, small finite values for n yield the best lower bounds.

5.4.3.4 A Lower Bound for Polynomial Total Edge Cost

Applying the lower bound on the price of stability

$$\text{PoS}(G_n \mid S_n) \geq \max_{n' \in [n]} \max_{k \in [n']} \text{PoS}(G_{n'}, c^{(z^k)} \mid S_{n'}, T_k)$$

for polynomial total edge cost gives similar results as for affine total edge cost. Again, choosing $k = \lceil n/2 \rceil$ provides the best bounds α large enough. This is illustrated by the diagrams in Figure 5.19. After maximizing over all $n \in \mathbb{N}$, we obtain the bound as depicted in Figure 5.20. Experiments, however, show that the lower bound obtained




Figure 5.20: Lower bounds to the price of stability for polynomial total edge cost. The bounds $\text{PoS}(G_n, c^{(z_{\lceil n/2 \rceil})} \mid S_n, T_{\lceil n/2 \rceil})$ are plotted for $n = 1$ (yellow) to $n = 100$ (violet). The bound resulting from maximizing over all $n \in N$ is highlighted in green. The upper bound based on Theorem 5.24 is plotted in red.

from the fan graph is not best possible for polynomial total edge cost. Table 5.3 shows the structure of worst case examples for small n under the assumption that the star is a Nash equilibrium that minimizes the social cost. These networks were determined by full enumeration. For large enough n and α , these worst case examples are not captured by the fan graph. Hence, further investigation should lead to improved lower bounds on the price of stability in this case.

5.5 Closing Remarks

In this chapter, we examine Nash equilibria in network cost-sharing games in terms of their computational complexity and their efficiency. Our analysis moves towards a better comprehension of cost functions between two extreme cases, constant total edge cost and linear total edge cost. The complexity of finding a Nash equilibrium does not seem to vary gradually in between. We obtain hardness for cost functions that are not linear. This vastly extends to finding special equilibria in multicast and broadcast games. The analysis of the running time for the improving dynamics in multicast and broadcast




Table 5.3: Trees T maximizing the price of stability under the assumption that the star S_n is a minimum-cost Nash equilibrium and T is a social optimum for $f(k) = k^{\alpha-1}$.

games seems to involve numerical intricacies. It remains open whether they converge in polynomial time.

The efficiency of Nash equilibria on the other hand appears to depend on the choice of the cost function in a more continuous way. A first evidence is given by the potential-based bounds on the price of anarchy and the price of stability. There, the study of affine and polynomial total edge costs gives smooth interpolations between the extremes. The same applies to the lower bounds on the price of stability that we obtain for these classes in broadcast games. Together, these findings suggest that constant and linear total edge costs yield the worst and best quality of equilibria, respectively. Intuitively, this can be explained with the benefit from sharing edges, which is larger the slower the total edge

cost grows. This effect can be leveraged by a social optimum, whereas it cannot by a Nash equilibrium due to lack of coordination. Formalizing this monotonicity and the involved properties of the cost function remains open. The upper bounds on the price of stability of broadcast games obtained from the homogenization-absorption framework, however, disagrees with this trend. Out of the class of functions that it applies to, constant total edge cost yields the best bounds. The reason for this lies in the arbitrary way that the homogenization and the absorption use edges in the support of a social optimum. High congestion of edges in the optimum is not utilized in the bounds. Improving on this promises better bounds on the price of stability and extended applicability of the method to a larger class of cost functions.

Notation

$ \cdot $	cardinality of a set. 5
$ \cdot $	value of a static/dynamic flow. 11
$[\cdot]$	first positive natural numbers. 5
$[\cdot]_-$	negative part of a number. 5
$[\cdot]_+$	positive part of a number. 5
$\ \cdot\ _\infty$	uniform norm of a vector or function. 6
\circ	composition of two relations or functions. 5
\cup	disjoint union of two sets. 5
\parallel	parallel composition of two two-terminal graphs. 11
$*$	series composition of two two-terminal graphs. 11
Δ	symmetric difference of two sets. 5
$\Rightarrow_{\mathcal{A}}$	disjoint paths relation. 27
$\Leftrightarrow_{\mathcal{A}_1, \mathcal{A}_2}$	two disjoint paths relation. 33
$\xrightarrow{\ell}_E$	disjoint shortest paths relation. 28
$\mathbb{1}$	all-one vector. 5
$\mathbb{1}_i$	unit vector with respect to component i . 5
$\mathbb{1}_X$	characteristic vector of a set X . 5
2^X	power set of a set X . 5
$A_T(v)$	ancestors of a vertex v in a rooted tree T . 10
$a_T^W(v)$	lowest ancestor of a vertex v from a set W in a rooted tree T . 10
$c_\sigma(e)$	cost for using edge e in strategy profile σ . 76
$C_i(\sigma)$	total cost of player i in strategy profile σ . 13, 76
$C(\sigma)$	social cost of strategy profile σ . 13, 77
$\delta_{\mathcal{A}}(U)$	arcs/edges in \mathcal{A} across the cut U . 7, 8
$\delta_A^-(U)$	incoming arcs of U in A . 8
$\delta_A^+(U)$	outgoing arcs of U in A . 8

Notation

$\vec{d}_{T,g}(v, w)$	directional broadcast distance from v to w in a tree T with respect to a cost function g . 118
$d_{T,g}(v, w)$	undirectional broadcast distance between v and w in a tree T with respect to a cost function g . 118
$D_T(v)$	descendants of a vertex v in a rooted tree T . 10
$D_T^W(u)$	directed descendants of a vertex w in a rooted tree T with respect to a set W . 10
$e_T(v)$	parent edge of a vertex v in a rooted tree T . 10
G/U	contraction of a set U in a graph G . 9
G/\mathcal{U}	contraction of a family of sets \mathcal{U} in a graph G . 9
$G[U]$	subgraph of a graph G induced by a set U . 9
\vec{G}	partial orientation of a graph G . 33
$H(n)$	n -th harmonic number. 5
Id	identity matrix. 5
$L_{\text{loc}}^1(\mathbb{R})$	locally Lebesgue-integrable functions on \mathbb{R} . 6
$\text{lca}_T(v, w)$	lowest common ancestor of vertices v and w in a rooted tree T . 10
$M_{I,J}$	submatrix of a matrix M with respect to rows I and columns J . 5
$M_{I,\bullet}$	submatrix of a matrix M with respect to rows I . 6
$M_{\bullet,J}$	submatrix of a matrix M with respect to columns J . 6
M^\top	transposed of a matrix M . 5
\mathbb{N}	natural numbers. 5
$n_\sigma(e)$	congestion of an edge e in a strategy profile σ . 76
$N_{\mathcal{A}}(v)$	neighbors of a vertex v with respect to a set of arcs/edges \mathcal{A} . 7, 8
$N_A^-(v)$	in-neighbors of a vertex v with respect to a set of arcs A . 8
$N_A^+(v)$	out-neighbors of a vertex v with respect to a set of arcs A . 8
$\mathcal{O}(f)$	set of functions not growing faster than f . 15
$\Omega(f)$	set of functions not growing slower than f . 16
$\pi(\ell')$	normalization of labels ℓ' of a thin flow with resetting. 52
PoS	price of stability. 131
PoS(G)	price of stability on a graph G . 131

$\text{PoS}(G \mid S)$	price of stability on a graph G with a cost-minimal Nash equilibrium S . 131
$\text{PoS}(G \mid S, T)$	largest ratio of social costs of S and T on a graph G with a cost-minimal Nash equilibrium S . 131
$\text{PoS}(G, c \mid S, T)$	ratio of social costs of S and T in a graph G with costs c . 131
\mathbb{R}	real numbers. 5
$\mathbb{R}_{\leq 0}$	nonpositive real numbers. 5
$\mathbb{R}_{< 0}$	negative real numbers. 5
$\mathbb{R}_{\geq 0}$	nonnegative real numbers. 5
$\mathbb{R}_{> 0}$	positive real numbers. 5
ϱ^a	label function on an arc a . 49
$\varrho^{G'}$	label function on a graph G' . 58
$(\sigma_{-i}, \hat{\sigma}_i)$	unilateral deviation of a player i in a strategy profile σ to $\hat{\sigma}_i$. 13
Σ_i	set of strategies of a player i . 13
Σ	set of strategy profiles. 13, 76
$T[v, w]$	unique v - w path in a tree T . 9
$\binom{V}{k}$	family of subsets of size k in a set V . 5
$x(I)$	sum of components I of a vector x . 6
$\chi^{G'}$	thin flow function on a graph G' . 58
x_I	subvector of a vector x with respect to indices I . 6
$\Phi(\sigma)$	potential of a strategy profile σ . 77

Index

- α -approximation, 17
- α -extension, 42, 64, 68
- absorption, 81, 112, 123
- ancestor, 10
 - lowest, 10
 - lowest common, 10
- arborescence, 10
- arc, 7
 - active, 46, 48
 - incoming, 8
 - outgoing, 8
 - resetting, 46, 48
- Bellman equations, 46, 49
- best response, 13
- capacity constraints, 11, 12, 44
- complementarity condition, 21, 49
- composition, 28
 - parallel, 11, 71
 - series, 11, 70
- congestion, 1, 76
- connected, 7, 8
 - strongly, 8
 - weakly, 8
- connected components, 7
 - strongly, 8
 - weakly, 8, 55, 58
- contraction, 9, 28, 55, 85
- covering, 16
- cut, 7
- cycle, 7
 - directed, 8
 - simple, 7
- depth, 10
- descendant, 10
 - direct, 10
- deviation
 - unilateral, 13, 77
 - improving, 13
- duality
 - strong, 19
 - weak, 19, 121
- edge, 7, 19
- equilibrium
 - α -approximate Nash, 15
 - dynamic, 40, 46
 - mixed Nash, 14
 - Nash, 13, 84
 - strong Nash, 15, 83
 - user, 14, 46
- excess, 11, 12
- exponential-time hypothesis, 18
- flow, 11
 - s - t , 11
 - decomposition, 12
 - multicommodity, 12, 27
 - over time, 12
 - path, 11
 - singlecommodity, 11
 - value, 11
- flow conservation, 49
 - strict, 11, 12, 44
 - weak, 11, 12
- flow over time, 40, 43
 - s - t , 12
 - feasible, 44
- fluid queuing model, 40, 43
- game

INDEX

- broadcast, 78, 87, 100, 131
- multicast, 77, 87, 104
- network cost-sharing, 75, 76
- normal form, 13
- graph, 7
 - directed, 8, 43, 77
 - directed acyclic, 10, 48
 - fan, 132
 - mixed, 8
 - planar, 10
 - series-parallel, 10, 43, 70, 73
 - two-terminal, 10
 - undirected, 7, 77
 - weakly acyclic, 27, 30
- harmonic mean
 - weighted, 6, 141, 144
- harmonic number, 5
- height, 10
- homogenization, 81, 112
- homogenization-absorption framework,
81, 112
- improving dynamics, 13, 89, 103
 - tree, 91
- incidence matrix, 9, 51
- incident, 7
- induced subgraph, 9
- instance, 15
- Laplacian matrix
 - weighted, 9, 54
- Lemke's algorithm, 21, 56
- linear complementarity problem, 21, 49
- linear program, 18, 120, 138
 - dual, 19
 - value, 18
- mediant, 6, 140, 145
 - inequality, 6
- monotone, 64
 - left-, 64, 65
 - right-, 43, 64, 65
- neighbor, 7, 8
 - in-, 8, 46, 49
 - out-, 8
- neighborhood, 17
 - polyhedral, 92
 - unilateral, 91
- network cost-sharing game
 - nonuniform, 78, 111
 - uniform, 78
- NP, 16
 - complete, 16
 - hard, 16, 101, 102
- optimum, 18, 102
 - local, 17, 91
- P, 16
- packing, 16
- partial orientation, 33
- path, 7, 8
 - directed, 8
 - disjoint, 25, 32
 - arc/edge-, 9, 24, 30, 32
 - internally vertex-, 9, 87
 - vertex-, 9, 24, 38
 - heavy, 97
 - shortest, 25, 32, 99
 - dynamic, 46, 48
 - simple, 7
- player, 12
- PLS, 17, 91
 - complete, 17, 95
- polyhedron, 19
- polytope, 19
- potential function, 77, 102, 106, 114
 - method, 80, 108
- PPAD, 17, 70
- price
 - of anarchy, 14, 42, 105, 106
 - of stability, 14, 80, 108, 126, 140
- problem, 15
 - disjoint paths, 24
 - disjoint shortest paths, 24
 - exact 3-set cover, 16, 100, 102
 - maximum cut, 17, 95

- reachable, 8
- reduction
 - of a decision problem, 16
 - of a local search problem, 17
- shortest path graph, 32
 - with resetting, 48
- social
 - cost, 13, 77
 - optimum, 13, 88, 132
- solution
 - basic, 18, 139
 - feasible, 16, 18
 - optimal, 18
- strategy, 13, 76
 - dominant, 13
 - profile, 13, 76
 - homogeneous, 117, 119
 - induced, 86
- thin flow with resetting, 41
 - normalized, 42, 49, 52, 60
 - parametric, 57
- time horizon, 12
- topological ordering, 10, 28, 52
- total edge cost, 76
 - affine, 78, 110, 136, 146
 - constant, 78, 136, 141
 - linear, 78
 - polynomial, 78, 110, 136, 150
- transit time, 12
- tree, 9, 86
 - rooted, 10
 - spanning, 9, 86
 - sub-, 10
- vertex, 7, 19

Bibliography

- Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. (1993). Network flows.
- Albers, S. (2008). [On the value of coordination in network design](#). S. Teng (ed.). *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008)*. SIAM, 294–303.
- Albers, S. (2009). [On the value of coordination in network design](#). *SIAM Journal on Computing* 38(6):2273–2302.
- Albers, S. and Lenzner, P. (2010). [On approximate Nash equilibria in network design](#). A. Saberi (ed.). *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE 2010)*. LNCS 6484. Springer, 14–25.
- Albers, S. and Lenzner, P. (2013). [On approximate Nash equilibria in network design](#). *Internet Mathematics* 9(4):384–405.
- Anderson, E. J. and Philpott, A. B. (1994). “Optimisation of flows in networks over time”. *Probability, Statistics and Optimisation*. Ed. by F. P. Kelly. Wiley, 369–382.
- Andrews, M. (2004). [Hardness of buy-at-bulk network design](#). *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS 2004)*. IEEE, 115–124.
- Anshelevich, E., Dasgupta, A., Kleinberg, J. M., Tardos, É., Wexler, T., and Roughgarden, T. (2004). [The price of stability for network design with fair cost allocation](#). *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS 2004)*. IEEE, 295–304.
- Anshelevich, E., Dasgupta, A., Kleinberg, J. M., Tardos, É., Wexler, T., and Roughgarden, T. (2008). [The price of stability for network design with fair cost allocation](#). *SIAM Journal on Computing* 38(4):1602–1623.
- Awerbuch, B. and Azar, Y. (1997). [Buy-at-bulk network design](#). *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS 1997)*. IEEE, 542–547.
- Bartal, Y. (1998). [On approximating arbitrary metrics by tree metrics](#). J. S. Vitter (ed.). *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC 1998)*. ACM, 161–168.
- Bentert, M., Nichterlein, A., Renken, M., and Zschoche, P. (2021). [Using a geometric lens to find k disjoint shortest paths](#). N. Bansal, E. Merelli, and J. Worrell (eds.). *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming*

Bibliography

- (ICALP 2021). LIPIcs 198. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 26:1–26:14.
- Bérczi, K. and Kobayashi, Y. (2017). [The directed disjoint shortest paths problem](#). K. Pruhs and C. Sohler (eds.). *Proceedings of the 25th Annual European Symposium on Algorithms (ESA 2017)*. LIPIcs 87. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 13:1–13:13.
- Bhaskar, U., Fleischer, L., and Anshelevich, E. (2011). [A stackelberg strategy for routing flow over time](#). D. Randall (ed.). *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)*. SIAM, 192–201.
- Bhaskar, U., Fleischer, L., and Anshelevich, E. (2015). [A stackelberg strategy for routing flow over time](#). *Games and Economic Behavior* 92:232–247.
- Bilò, V. and Bove, R. (2011). [Bounds on the price of stability of undirected network design games with three players](#). *Journal of Interconnection Networks* 12(1-2):1–17.
- Bilò, V., Caragiannis, I., Fanelli, A., and Monaco, G. (2010). [Improved lower bounds on the price of stability of undirected network design games](#). S. C. Kontogiannis, E. Koutsoupias, and P. G. Spirakis (eds.). *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT 2010)*. LNCS 6386. Springer, 90–101.
- Bilò, V., Caragiannis, I., Fanelli, A., and Monaco, G. (2013). [Improved lower bounds on the price of stability of undirected network design games](#). *Theory of Computing Systems* 52(4):668–686.
- Bilò, V., Flammini, M., Monaco, G., and Moscardelli, L. (2015). [Computing approximate Nash equilibria in network congestion games with polynomially decreasing cost functions](#). E. Markakis and G. Schäfer (eds.). *Proceedings of the 11th International Conference on Web and Internet Economics (WINE 2015)*. LNCS 9470. Springer, 118–131.
- Bilò, V., Flammini, M., Monaco, G., and Moscardelli, L. (2021). [Computing approximate Nash equilibria in network congestion games with polynomially decreasing cost functions](#). *Distributed Computing* 34(1):1–14.
- Bilò, V., Flammini, M., and Moscardelli, L. (2013). [The price of stability for undirected broadcast network design with fair cost allocation is constant](#). *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*. IEEE, 638–647.
- Bilò, V., Flammini, M., and Moscardelli, L. (2020). [The price of stability for undirected broadcast network design with fair cost allocation is constant](#). *Games and Economic Behavior* 123:359–376.
- Björklund, A. and Husfeldt, T. (2014). [Shortest two disjoint paths in polynomial time](#). J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias (eds.). *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP 2014), Part I*. LNCS 8572. Springer, 211–222.

- Björklund, A. and Husfeldt, T. (2019). [Shortest two disjoint paths in polynomial time](#). *SIAM Journal on Computing* 48(6):1698–1710.
- Borwein, D. and Borwein, J. M. (1995). [On an intriguing integral and some series related to \$\zeta\(4\)\$](#) . *Proceedings of the American Mathematical Society* 123(4):1191–1198.
- Braess, D. (1968). [Über ein Paradoxon aus der Verkehrsplanung](#). German. *Unternehmensforschung* 12(1):258–268. Translated as: D. Braess, A. Nagurney, and T. Wakolbinger. [On a paradox of traffic planning](#). *Transportation Science* 39(4) (2005):446–450.
- Brandstädt, A., Le, V., and Spinrad, J. (1999). [Graph Classes: A Survey](#). SIAM.
- Brouwer, L. E. J. (1911). [Über Abbildung von Mannigfaltigkeiten](#). German. *Mathematische Annalen* 71 (1):97–115.
- Burkard, R. E., Dlaska, K., and Klinz, B. (1993). [The quickest flow problem](#). *Mathematical Methods of Operations Research* 37(1):31–58.
- Charikar, M., Karloff, H., Mathieu, C., Naor, J. S., and Saks, M. (2008). [Online multicast with egalitarian cost sharing](#). F. M. auf der Heide and N. Shavit (eds.). *Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2008)*. ACM, 70–76.
- Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., and Orda, A. (2006). [Non-cooperative multicast and facility location games](#). J. Feigenbaum, J. C. Chuang, and D. M. Pennock (eds.). *Proceedings of the 7th ACM Conference on Electronic Commerce (EC 2006)*. ACM, 72–81.
- Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., and Orda, A. (2007). [Non-cooperative multicast and facility location games](#). *IEEE Journal on Selected Areas in Communications* 25(6):1193–1206.
- Chekuri, C., Khanna, S., and Naor, J. S. (2001). [A deterministic algorithm for the cost-distance problem](#). S. R. Kosaraju (ed.). *Proceedings of the 12th Annual Symposium on Discrete Algorithms (SODA 2001)*. ACM/SIAM, 232–233.
- Chen, H. and Roughgarden, T. (2009). [Network design with weighted players](#). *Theory of Computing Systems* 45(2):302–324.
- Chen, H.-L. and Roughgarden, T. (2006). [Network design with weighted players](#). P. B. Gibbons and U. Vishkin (eds.). *Proceedings of the 18th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2006)*. ACM, 29–38.
- Chen, X., Deng, X., and Teng, S.-H. (2009). [Settling the complexity of computing two-player Nash equilibria](#). *Journal of the ACM* 56(3).
- Christodoulou, G., Chung, C., Ligett, K., Pyrga, E., and van Stee, R. (2009). [On the price of stability for undirected network design](#). E. Bampis and K. Jansen (eds.). *Proceedings of the 7th International Workshop on Approximation and Online Algorithms (WAOA 2009)*. LNCS 5893. Springer, 86–97.

Bibliography

- Chuzhoy, J., Gupta, A., Naor, J., and Sinha, A. (2005). [On the approximability of some network design problems](#). *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*. SIAM, 943–951.
- Chuzhoy, J., Gupta, A., Naor, J., and Sinha, A. (2008). [On the approximability of some network design problems](#). *ACM Transactions on Algorithms* 4(2):23:1–23:17.
- Cominetti, R., Correa, J. R., and Larré, O. (2011). [Existence and uniqueness of equilibria for flows over time](#). L. Aceto, M. Henzinger, and J. Sgall (eds.). *Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP 2011), Part II*. LNCS 6756. Springer, 552–563.
- Cominetti, R., Correa, J. R., and Larré, O. (2015). [Dynamic equilibria in fluid queueing networks](#). *Operations Research* 63(1):21–34.
- Cominetti, R., Correa, J. R., and Olver, N. (2017). [Long term behavior of dynamic equilibria in fluid queueing networks](#). F. Eisenbrand and J. Könemann (eds.). *Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization (IPCO 2017)*. LNCS 10328. Springer, 161–172.
- Cominetti, R., Harks, T., Osorio, C., and Peis, B. (2018). [Dynamic traffic models in transportation science \(Dagstuhl Seminar 18102\)](#). *Dagstuhl Reports* 8(3):21–38.
- Cook, S. A. (1971). [The complexity of theorem-proving procedures](#). M. A. Harrison, R. B. Banerji, and J. D. Ullman (eds.). *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing (STOC 1971)*. ACM, 151–158.
- Correa, J. R., Cristi, A., and Oosterwijk, T. (2019). [On the price of anarchy for flows over time](#). A. Karlin, N. Immorlica, and R. Johari (eds.). *Proceedings of the 20th ACM Conference on Economics and Computation (EC 2019)*. ACM, 559–577.
- Cottle, R. W., Pang, J.-S., and Stone, R. E. (2009). [The Linear Complementarity Problem](#). SIAM.
- Dantzig, G. B. (1990). [Origins of the simplex method](#). *A History of Scientific Computing*. ACM, 141–151.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2006). [The complexity of computing a Nash equilibrium](#). J. M. Kleinberg (ed.). *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006)*. ACM, 71–78.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). [The complexity of computing a nash equilibrium](#). *SIAM Journal on Computing* 39(1):195–259.
- Dinic, E. A. (1970). Algorithm for solution of a problem of maximum flow in a network with power estimation. *Soviet Mathematics - Doklady* 11:1277–1280.
- Disser, Y., Feldmann, A. E., Klimm, M., and Mihalák, M. (2013). [Improving the \$H\(k\)\$ -bound on the price of stability in undirected Shapley network design games](#). P. G. Spirakis and M. J. Serna (eds.). *Proceedings of the 8th International Conference on Algorithms and Complexity (CIAC 2013)*. LNCS 7878. Springer, 158–169.

- Disser, Y., Feldmann, A. E., Klimm, M., and Mihalák, M. (2015). [Improving the \$H\(k\)\$ -bound on the price of stability in undirected Shapley network design games](#). *Theoretical Computer Science* 562:557–564.
- Edmonds, J. and Karp, R. M. (1972). [Theoretical improvements in algorithmic efficiency for network flow problems](#). *Journal of the ACM* 19(2):248–264.
- Eilam-Tzoref, T. (1998). [The disjoint shortest paths problem](#). *Discrete Applied Mathematics* 85(2):113–138.
- Elias, P., Feinstein, A., and Shannon, C. (1956). [A note on the maximum flow through a network](#). *IRE Transactions on Information Theory* 2(4):117–119.
- Epstein, A., Feldman, M., and Mansour, Y. (2007). [Strong equilibrium in cost sharing connection games](#). J. K. MacKie-Mason, D. C. Parkes, and P. Resnick (eds.). *Proceedings of the 8th ACM Conference on Electronic Commerce (EC 2007)*. ACM, 84–92.
- Epstein, A., Feldman, M., and Mansour, Y. (2009). [Strong equilibrium in cost sharing connection games](#). *Games and Economic Behavior* 67(1):51–68.
- Erickson, R. E., Monma, C. L., and Veinott, A. F. (1987). [Send-and-split method for minimum-concave-cost network flows](#). *Mathematics of Operations Research* 12(4):634–664.
- Even, S., Itai, A., and Shamir, A. (1975). [On the complexity of timetable and multi-commodity flow problems](#). *Proceedings of the 16th Annual Symposium on Foundations of Computer Science (FOCS 1975)*. IEEE, 184–193.
- Even, S., Itai, A., and Shamir, A. (1976). [On the complexity of timetable and multicommodity flow problems](#). *SIAM Journal on Computing* 5(4):691–703.
- Fakcharoenphol, J., Rao, S., and Talwar, K. (2003). [A tight bound on approximating arbitrary metrics by tree metrics](#). L. L. Larmore and M. X. Goemans (eds.). *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC 2003)*. ACM, 448–455.
- Fakcharoenphol, J., Rao, S., and Talwar, K. (2004). [A tight bound on approximating arbitrary metrics by tree metrics](#). *Journal of Computer and System Sciences* 69(3):485–497.
- Fanelli, A., Leniowski, D., Monaco, G., and Sankowski, P. (2012). [The ring design game with fair cost allocation](#). P. W. Goldberg (ed.). *Proceedings of the 8th International Workshop on Internet and Network Economics (WINE 2012)*. LNCS 7695. Springer, 546–552.
- Fanelli, A., Leniowski, D., Monaco, G., and Sankowski, P. (2015). [The ring design game with fair cost allocation](#). *Theoretical Computer Science* 562:90–100.
- Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., and Shabo, R. (2006). [On the price of stability for designing undirected networks with fair cost allocations](#). M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener (eds.). *Proceedings of the 33rd International Collo-*

Bibliography

- quium on Automata, Languages and Programming (ICALP 2006), Part I*. LNCS 4051. Springer, 608–618.
- Filippov, A. F. (1988). [Differential Equations with Discontinuous Righthand Sides](#). Ed. by F. M. Arscott. Vol. 18. Mathematics and its Applications (Soviet Series). Translated from Russian. Springer.
- Fleischer, L. and Tardos, É. (1998). [Efficient continuous-time dynamic network flow algorithms](#). *Operations Research Letters* 23(3):71–80.
- Fleischer, L. and Skutella, M. (2007). [Quickest flows over time](#). *SIAM Journal on Computing* 36(6):1600–1630.
- Ford, L. R. and Fulkerson, D. R. (1958). [Constructing maximal dynamic flows from static flows](#). *Operations Research* 6(3):419–433.
- Ford Jr., L. R. and Fulkerson, D. R. (1956). [Maximal flow through a network](#). *Canadian Journal of Mathematics* 8:399–404.
- Ford, L. R. and Fulkerson, D. R. (1962). [Flows in Networks](#). Princeton University Press.
- Fortune, S., Hopcroft, J., and Wyllie, J. (1980). [The directed subgraph homeomorphism problem](#). *Theoretical Computer Science* 10(2):111–121.
- Freeman, R., Haney, S., and Panigrahi, D. (2016). [On the price of stability of undirected multicast games](#). Y. Cai and A. Vetta (eds.). *Proceedings of the 12th International Conference on Web and Internet Economics (WINE 2016)*. LNCS 10123. Springer, 354–368.
- Gale, D. (1959). [Transient flows in networks](#). *Michigan Mathematical Journal* 6(1):59–63.
- Gallo, G. and Sodini, C. (1979). [Adjacent extreme flows and application to min concave cost flow problems](#). *Networks* 9(2):95–121.
- Garey, M. R. and Johnson, D. S. (1990). *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.
- Geršgorin, S. (1931). Über die Abgrenzung der Eigenwerte einer Matrix. German. *Izvestiya Rossiiskoi Akademii Nauk, Seriya Matematicheskaya* 6:749–754.
- Goldberg, P. W. and Papadimitriou, C. H. (2018a). [Towards a unified complexity theory of total functions](#). A. R. Karlin (ed.). *Proceedings of the 9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Vol. 94. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 37:1–37:20.
- Goldberg, P. W. and Papadimitriou, C. H. (2018b). [Towards a unified complexity theory of total functions](#). *Journal of Computer and System Sciences* 94:167–192.
- Gottschau, M., Kaiser, M., and Waldmann, C. (2019). [The undirected two disjoint shortest paths problem](#). *Operations Research Letters* 47(1):70–75.
- Graf, L. and Harks, T. (2019). [Dynamic flows with adaptive route choice](#). A. Lodi and V. Nagarajan (eds.). *Proceedings of the 20th International Conference on Integer Pro-*

- gramming and Combinatorial Optimization (IPCO 2019)*. LNCS 11480. Springer, 219–232.
- Graf, L., Harks, T., and Sering, L. (2020). [Dynamic flows with adaptive route choice](#). *Mathematical Programming* 183(1):309–335.
- Guha, S., Meyerson, A., and Munagala, K. (2001). [A constant factor approximation for the single sink edge installation problems](#). J. S. Vitter, P. G. Spirakis, and M. Yannakakis (eds.). *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*. ACM, 383–388.
- Guisewite, G. M. and Pardalos, P. M. (1993). [A polynomial time solvable concave network flow problem](#). *Networks* 23(2):143–147.
- Guisewite, G. M. and Pardalos, P. M. (1991). [Algorithms for the single-source uncapacitated minimum concave-cost network flow problem](#). *Journal of Global Optimization* 1(3):245–265.
- Gupta, A. and Könemann, J. (2011). [Approximation algorithms for network design: a survey](#). *Surveys in Operations Research and Management Science* 16(1):3–20.
- Gupta, A., Kumar, A., Pál, M., and Roughgarden, T. (2007). [Approximation via cost sharing: simpler and better approximation algorithms for network design](#). *Journal of the ACM* 54(3):11:1–11:38.
- Gupta, A., Kumar, A., and Roughgarden, T. (2003). [Simpler and better approximation algorithms for network design](#). L. L. Larmore and M. X. Goemans (eds.). *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC 2003)*. ACM, 365–372.
- Hardy, G., Littlewood, J., and Pólya, G. (1953). [Inequalities](#). *The Mathematical Gazette* 37(321):236–236.
- Hassani, M. (2005). [Approximation of the Lambert W function](#). *RGMIA Research report collection* 8(4).
- Hoppe, B. and Tardos, É. (1994). [Polynomial time algorithms for some evacuation problems](#). D. D. Sleator (ed.). *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1994)*. ACM/SIAM, 433–441.
- Hoppe, B. and Tardos, É. (1995). [The quickest transshipment problem](#). K. L. Clarkson (ed.). *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1995)*. ACM/SIAM, 512–521.
- Hoppe, B. and Tardos, É. (2000). [The quickest transshipment problem](#). *Mathematics of Operations Research* 25(1):36–62.
- Impagliazzo, R. and Paturi, R. (1999). [Complexity of \$k\$ -sat](#). D. C. Martin (ed.). *Proceedings of the 14th Annual IEEE Conference on Computational Complexity*. IEEE Computer Society, 237–240.

Bibliography

- Impagliazzo, R. and Paturi, R. (2001). [On the complexity of \$k\$ -sat](#). *Journal of Computer and System Sciences* 62(2):367–375.
- Israel, J. and Sering, L. (2020). [The impact of spillback on the price of anarchy for flows over time](#). T. Harks and M. Klimm (eds.). *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT 2020)*. LNCS 12283. Springer, 114–129.
- Jaffe, A. M. (2006). [The millennium grand challenge in mathematics](#). *Notices of the AMS* 53(6):652–660.
- Johnson, D. S., Papadimitriou, C. H., and Yannakakis, M. (1985). [How easy is local search?](#) *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS 1985)*. IEEE, 39–42.
- Johnson, D. S., Papadimitriou, C. H., and Yannakakis, M. (1988). [How easy is local search?](#) *Journal of Computer and System Sciences* 37(1):79–100.
- Kahn, A. B. (1962). [Topological sorting of large networks](#). *Communications of the ACM* 5(11):558–562.
- Kaiser, M. (2022). [Computation of dynamic equilibria in series-parallel networks](#). *Mathematics of Operations Research* 47(1):50–71.
- Kakutani, S. (1941). [A generalization of Brouwer’s fixed point theorem](#). *Duke Mathematical Journal* 8(3):457–459.
- Kantorovich, L. V. (1939). *Matematicheskiye metody organizatsii i planirovaniya proizvodstva*. Russian. Leningrad State University. Translated as: [Mathematical Methods of Organizing and Planning Production](#). Trans. by W. Campbell and W. H. Marlow. *Management Science* 6(4) (1960):366–422.
- Karp, R. M. (1972). [Reducibility among combinatorial problems](#). R. E. Miller, J. W. Thatcher, and J. D. Bohlinger (eds.). *Proceedings of a Symposium on the Complexity of Computer Computations*. Springer, 85–103.
- Karp, R. M. (1975). [On the computational complexity of combinatorial problems](#). *Networks* 5(1):45–68.
- Kawarabayashi, K.-i., Kobayashi, Y., and Reed, B. (2012). [The disjoint paths problem in quadratic time](#). *Journal of Combinatorial Theory, Series B* 102(2):424–435.
- Kawase, Y. and Makino, K. (2012). [Nash equilibria with minimum potential in undirected broadcast games](#). M. S. Rahman and S. Nakano (eds.). *Proceedings of the 6th International Workshop on Algorithms and Computation (WALCOM 2012)*. LNCS 7157. Springer, 217–228.
- Kawase, Y. and Makino, K. (2013). [Nash equilibria with minimum potential in undirected broadcast games](#). *Theoretical Computer Science* 482:33–47.
- Khachiyan, L. G. (1979). A polynomial algorithm in linear programming. Russian. *Proceedings of the USSR Academy of Sciences* 244(5):1093–1096.

- Klee, V. and Minty, G. J. (1972). How good is the simplex algorithm? O. Shisha (ed.). *Inequalities – III*. Academic Press, 159–175.
- Klinz, B. and Tuy, H. (1993). [Minimum concave-cost network flow problems with a single nonlinear arc cost](#). D. Z. Du and P. M. Pardalos (eds.). *Network Optimization Problems: Algorithms, Applications and Complexity*. World Scientific, 125–145.
- Klinz, B. and Woeginger, G. J. (1995). [Minimum cost dynamic flows: the series-parallel case](#). E. Balas and J. Clausen (eds.). *Proceedings of the 4th International Conference on Integer Programming and Combinatorial Optimization (IPCO 1995)*. LNCS 920. Springer, 329–343.
- Klinz, B. and Woeginger, G. J. (2004). [Minimum-cost dynamic flows: the series-parallel case](#). *Networks* 43(3):153–162.
- Kobayashi, Y. and Sako, R. (2019). [Two disjoint shortest paths problem with non-negative edge length](#). *Operations Research Letters* 47(1):66–69.
- Koch, R. and Skutella, M. (2009). [Nash equilibria and the price of anarchy for flows over time](#). M. Mavronicolas and V. G. Papadopoulou (eds.). *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT 2009)*. LNCS 5814. Springer, 323–334.
- Koch, R. and Skutella, M. (2011). [Nash equilibria and the price of anarchy for flows over time](#). *Theory of Computing Systems* 49(1):71–97.
- Kollock, P. (1998). [Social dilemmas: the anatomy of cooperation](#). *Annual Review of Sociology* 24(1):183–214.
- Kotnyek, B. (2003). *An annotated overview of dynamic network flows*. Tech. rep. RR-4936. INRIA.
- Koutsoupias, E. and Papadimitriou, C. (1999). [Worst-case equilibria](#). C. Meinel and S. Tison (eds.). *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS 1999)*. LNCS 1563. Springer, 404–413.
- Koutsoupias, E. and Papadimitriou, C. (2009). [Worst-case equilibria](#). *Computer Science Review* 3(2):65–69.
- Kuratowski, C. (1922). [Une méthode d’élimination des nombres transfinis des raisonnements mathématiques](#). French. *Fundamenta Mathematicae* 3(1):76–108.
- Lee, E. and Ligett, K. (2013). [Improved bounds on the price of stability in network cost sharing games](#). M. J. Kearns, R. P. McAfee, and É. Tardos (eds.). *Proceedings of the 14th ACM Conference on Electronic Commerce (EC 2013)*. ACM, 607–620.
- Lemke, C. E. (1965). [Bimatrix equilibrium points and mathematical programming](#). *Management Science* 11(7):681–689.
- Lemke, C. E. and Howson Jr., J. T. (1964). [Equilibrium points of bimatrix games](#). *Journal of the Society for Industrial and Applied Mathematics* 12(2):413–423.

Bibliography

- Li, J. (2009). [An \$\mathcal{O}\(\log\(n\)/\log\(\log\(n\)\)\)\$ upper bound on the price of stability for undirected Shapley network design games.](#) *Information Processing Letters* 109(15):876–878.
- Lochet, W. (2021). [A polynomial time algorithm for the \$k\$ -disjoint shortest paths problem.](#) D. Marx (ed.). *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*. SIAM, 169–178.
- Lund, C. and Yannakakis, M. (1993). [On the hardness of approximating minimization problems.](#) S. R. Kosaraju, D. S. Johnson, and A. Aggarwal (eds.). *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC 1993)*. ACM, 286–293.
- Lund, C. and Yannakakis, M. (1994). [On the hardness of approximating minimization problems.](#) *Journal of the ACM* 41(5):960–981.
- Lynch, J. F. (1975). [The equivalence of theorem proving and the interconnection problem.](#) *ACM SIGDA Newsletter* 5(3):31–36.
- Macko, M., Larson, K., and Steskal, L. (2010). [Braess’s paradox for flows over time.](#) S. C. Kontogiannis, E. Koutsoupias, and P. G. Spirakis (eds.). *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT 2010)*. LNCS 6386. Springer, 262–275.
- Macko, M., Larson, K., and Steskal, L. (2013). [Braess’s paradox for flows over time.](#) *Theory Comput. Syst.* 53(1):86–106.
- Mamageishvili, A., Mihalák, M., and Montemezzani, S. (2018). [Improved bounds on equilibria solutions in the network design game.](#) *International Journal of Game Theory* 47(4):1113–1135.
- Mamageishvili, A., Mihalák, M., and Montemezzani, S. (2014). [An \$H\(n/2\)\$ upper bound on the price of stability of undirected network design games.](#) E. Csuhaaj-Varjú, M. Dietzfelbinger, and Z. Ésik (eds.). *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science (MFCS 2014)*. LNCS 8635. Springer, 541–552.
- Megiddo, N. (1978). [Combinatorial optimization with rational objective functions.](#) R. J. Lipton, W. A. Burkhard, W. J. Savitch, E. P. Friedman, and A. V. Aho (eds.). *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC 1978)*. ACM, 1–12.
- Megiddo, N. (1979). [Combinatorial optimization with rational objective functions.](#) *Mathematics of Operations Research* 4(4):414–424.
- Menger, K. (1927). [Zur allgemeinen Kurventheorie.](#) German. *Fundamenta Mathematicae* 10(1):96–115.
- Merchant, D. K. and Nemhauser, G. L. (1978a). [A model and an algorithm for the dynamic traffic assignment problems.](#) *Transportation Science* 12(3):183–199.
- Merchant, D. K. and Nemhauser, G. L. (1978b). [Optimality conditions for a dynamic traffic assignment model.](#) *Transportation Science* 12(3):200–207.

- Meunier, F. and Wagner, N. (2010). [Equilibrium results for dynamic congestion games](#). *Transportation Science* 44(4):524–536.
- Meyerson, A., Munagala, K., and Plotkin, S. A. (2000). [Cost-distance: two metric network design](#). *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*. IEEE, 624–630.
- Meyerson, A., Munagala, K., and Plotkin, S. A. (2008). [Cost-distance: two metric network design](#). *SIAM Journal on Computing* 38(4):1648–1659.
- Minieka, E. (1973). [Maximal, lexicographic, and dynamic network flows](#). *Operations Research* 21(2):517–527.
- Monderer, D. and Shapley, L. S. (1996). [Potential games](#). *Games and Economic Behavior* 14(1):124–143.
- Nash, J. F. (1950). [Equilibrium points in \$n\$ -person games](#). *Proceedings of the National Academy of Sciences of the USA* 36(1):48–49.
- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). [Algorithmic Game Theory](#). Cambridge University Press.
- Orlin, J. B., Punnen, A. P., and Schulz, A. S. (2009). [Integer programming: optimization and evaluation are equivalent](#). F. K. H. A. Dehne, M. L. Gavrilova, J. Sack, and C. D. Tóth (eds.). *Proceedings of the 11th International Symposium on Algorithms and Data Structures (WADS 2009)*. Vol. 5664. LNCS. Springer, 519–529.
- Papadimitriou, C. H. (1994). [On the complexity of the parity argument and other inefficient proofs of existence](#). *Journal of Computer and System Sciences* 48(3):498–532.
- Papadimitriou, C. H. (2001). [Algorithms, games, and the internet](#). F. Orejas, P. G. Spirakis, and J. van Leeuwen (eds.). *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*. LNCS 2076. Springer, 1–3.
- Peeta, S. and Ziliaskopoulos, A. K. (2001). [Foundations of dynamic traffic assignment: the past, the present and the future](#). *Networks and Spatial Economics* 1(3):233–265.
- Pham, H. M. and Sering, L. (2020). [Dynamic equilibria in time-varying networks](#). T. Harks and M. Klimm (eds.). *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT 2020)*. LNCS 12283. Springer, 130–145.
- Robertson, N. and Seymour, P. D. (1995). [Graph minors .XIII. The disjoint paths problem](#). *Journal of Combinatorial Theory, Series B* 63(1):65–110.
- Rosenthal, R. W. (1973). [A class of games possessing pure-strategy Nash equilibria](#). *International Journal of Game Theory* 2:65–67.
- Roughgarden, T. (2009). [Intrinsic robustness of the price of anarchy](#). M. Mitzenmacher (ed.). *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009)*. ACM, 513–522.
- Roughgarden, T. (2015). [Intrinsic robustness of the price of anarchy](#). *Journal of the ACM* 62(5):32:1–32:42.

Bibliography

- Saho, M. and Shigeno, M. (2017). [Cancel-and-tighten algorithm for quickest flow problems](#). *Networks* 69(2):179–188.
- Schäffer, A. A. and Yannakakis, M. (1991). [Simple local search problems that are hard to solve](#). *SIAM Journal on Computing* 20(1):56–87.
- Schlöter, M. and Skutella, M. (2017). [Fast and memory-efficient algorithms for evacuation problems](#). P. N. Klein (ed.). *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*. SIAM, 821–840.
- Schlöter, M., Skutella, M., and Van Tran, K. (2021). [A faster algorithm for quickest transshipments via an extended discrete newton method](#). *CoRR* abs/2108.06239.
- Schrijver, A. (1998). *Theory of linear and integer programming*. John Wiley & Sons.
- Schulz, A. S. (2009). [On the relative complexity of 15 problems related to 0/1-integer programming](#). W. J. Cook, L. Lovász, and J. Vygen (eds.). *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization*. Springer, 399–428.
- Schulz, A. S. and Stier-Moses, N. (2003). [On the performance of user equilibria in traffic networks](#). *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2003)*. SIAM, 86–87.
- Sering, L. and Skutella, M. (2018). [Multi-source multi-sink nash flows over time](#). R. Borndörfer and S. Storandt (eds.). *Proceedings of the 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)*. OASICS 65. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 12:1–12:20.
- Sering, L. and Vargas Koch, L. (2019). [Nash flows over time with spillback](#). T. M. Chan (ed.). *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*. SIAM, 935–945.
- Seymour, P. D. (1980). [Disjoint paths in graphs](#). *Discrete Mathematics* 29(3):293–309.
- Shapley, L. S. (1951). [Notes on the N-Person Game — II: The Value of an N-Person Game](#). RAND Corporation.
- Shiloach, Y. (1980). [A polynomial solution to the undirected two paths problem](#). *Journal of the ACM* 27(3):445–456.
- Skutella, M. (2009). [An introduction to network flows over time](#). W. J. Cook, L. Lovász, and J. Vygen (eds.). *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization*. Springer, 451–482.
- Spielman, D. A. and Teng, S. (2001). [Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time](#). J. S. Vitter, P. G. Spirakis, and M. Yannakakis (eds.). *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*. ACM, 296–305.
- Spielman, D. A. and Teng, S. (2004). [Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time](#). *Journal of the ACM* 51(3):385–463.

- Suurballe, J. W. (1974). [Disjoint paths in a network](#). *Networks* 4(2):125–145.
- Syrngkanis, V. (2010). [The complexity of equilibria in cost sharing games](#). A. Saberi (ed.). *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE 2010)*. LNCS 6484. Springer, 366–377.
- Talwar, K. (2002). [The single-sink buy-at-bulk LP has constant integrality gap](#). W. J. Cook and A. S. Schulz (eds.). *Proceedings of the 9th International Conference on Integer Programming and Combinatorial Optimization (IPCO 2002)*. LNCS 2337. Springer, 475–486.
- Thomassen, C. (1980). [2-linked graphs](#). *European Journal of Combinatorics* 1(4):371–378.
- Tuy, H. (2000). [The MCCNF problem with a fixed number of nonlinear arc costs: complexity and approximation](#). P. M. Pardalos (ed.). *Approximation and Complexity in Numerical Optimization: Continuous and Discrete Problems*. Springer, 525–544.
- Tuy, H., Ghannadan, S., Migdalas, A., and Värbrand, P. (1995). [The minimum concave cost network flow problem with fixed numbers of sources and nonlinear arc costs](#). *Journal of Global Optimization* 6(2):135–151.
- Vickrey, W. S. (1969). [Congestion theory and transport investment](#). *The American Economic Review* 59(2):251–260.
- Wardrop, J. G. (1952). [Some theoretical aspects of road traffic research](#). *Proceedings of the Institution of Civil Engineers* 1(3):325–362.
- Wilkinson, W. L. (1971). [An algorithm for universal maximal dynamic flows in a network](#). *Operations Research* 19(7):1602–1612.
- Yagar, S. (1971). [Dynamic traffic assignment by individual path minimization and queuing](#). *Transportation Research* 5(3):179–196.
- Zadeh, N. (1973). [A bad network problem for the simplex method and other minimum cost flow algorithms](#). *Mathematical Programming* 5(1):255–266.
- Zangwill, W. I. (1968). [Minimum concave cost flows in certain networks](#). *Management Science* 14(7):429–450.
- Zhang, C. and Nagamochi, H. (2012). [The next-to-shortest path in undirected graphs with nonnegative weights](#). J. Mestre (ed.). *Proceedings of the 18th Computing: The Australasian Theory Symposium (CATS 2012)*. CRPIT 128. Australian Computer Society, 13–20.
- Zhu, D. and Marcotte, P. (2000). [On the existence of solutions to the dynamic user equilibrium problem](#). *Transportation Science* 34(4):402–414.
- Zorn, M. (1935). [A remark on method in transfinite algebra](#). *Bulletin of the American Mathematical Society* 41(10):667–670.