Dissertation

Selfish Routing with Incomplete Information

Karsten Tiemann

Schriftliche Arbeit zur Erlangung des Grades
Doktor der Naturwissenschaften

Universität Paderborn,
Fakultät für Elektrotechnik, Informatik und Mathematik

Paderborn, 14. Juni 2007

Abstract

To study *selfish routing* scenarios in *networks* we use and extend in this thesis two well-known classes of games modeling such routing scenarios: network congestion games and Wardrop games. In both games, we are given a *network* with *edge latency functions*. In a *network congestion game*, each *player* selects as its *strategy* one path from its origin to its destination node and experiences as its *private cost* the sum of edge latencies on this path. In a *Nash equilibrium*, no player can decrease its private cost by unilaterally deviating to another path. In a *Wardrop game*, amounts of traffic are associated with pairs of network nodes. The traffic from an origin to a destination node is modeled as a splittable network flow and the cost on an origin-destination path is again given by the sum of edge latencies on this path. In a *Wardrop equilibrium*, no fraction of the traffic assigned to some path, however small, can decrease its cost by unilaterally switching to another path.

This thesis is primarily concerned with network routing scenarios where the players have *incomplete information*. One possibility to model such scenarios is to assume that a player who does not know some relevant parameter of the game is at least aware of a probability distribution over the possible outcomes of this parameter. In such a setting, it is reasonable to assume that the decisions of a player are based on the expected values of the unknown parameters. We apply this approach for network routing games where the players have incomplete information about the edge latency functions. Since each player obtains for each edge his own expected latency function we get games with *player-specific* latency functions. For both network congestion games and Wardrop games with player-specific latency functions, we show positive and negative results concerning the convergence to equilibria, the existence and polynomial-time computability of equilibria. We also prove bounds on the so-called *price of anarchy* that measures the worst-possible inefficiency of equilibria with respect to a social welfare measure.

We use an incomplete information model different from the aforementioned one for games where, in contrast to congestion games, the players do not know each other's weight. Based on Harsanyi's incomplete information concept of Bayesian games, each player in our *Bayesian routing games* has a set of possible types and each *type* of a player corresponds to some weight. The players' uncertainty about each other's weight is described by one *probability distribution* over all possible type profiles that is known to all players. In this setting, we focus on the price of anarchy, the existence and the computational complexity of equilibria.

We also study in this thesis, as a complete information setting, *bottleneck games with splittable traffic* where the latency on a path is given by the *maximum* latency of an edge on this path. We characterize for which games the social welfare of equilibria is unique and we give results on the *price of stability* that measures the worst-possible inefficiency of the best equilibrium.

Acknowledgments

My acknowledgments go to all persons who supported my research during the past three years. First and foremost, I would like to thank my advisor, Burkhard Monien, whose scientific expertise and constant guidance gave me an excellent support in the process of my scientific work. Moreover, I would like to express my thanks to Yvonne Bleischwitz, Martin Gairing, and Florian Schoppmann for the great collaboration on game theory topics. I am indebted to Martin Gairing, Florian Schoppmann, and Ulf-Peter Schroeder for the careful reading of parts of this dissertation and for numerous suggestions for improvements. I would also like to thank all members of the research group Monien for the very good cooperation and the excellent working atmosphere.

Thanks go to Marios Mavronicolas and Igal Milchtaich for inspiring my work with many good ideas and to Michael Dellnitz and Friedhelm Meyer auf der Heide for being my Ph.D. co-advisors. Furthermore, I would like to thank the International Graduate School of Dynamic Intelligent Systems at the University of Paderborn for establishing a stimulating environment for research and for providing financial support in the form of a scholarship.

Last but not least, I am very grateful to my family and friends whose continuous support greatly helped me to finish this thesis.

Paderborn, June 2007 Karsten Tiemann

Contents

CHAPTER 1

Introduction

1.1 Motivation and Framework

In this thesis, we study *selfish* behavior in *network routing games*. A selfish entity is concerned with its own interests but not with the interests of others. There is a bunch of different reasons why entities behave in a selfish way. Some systems have to rely on local selfish decisions of agents since it is impossible to set up a central control unit that supervises the system. This applies to many systems that are too large or too dynamic for a central control unit. For example it is impossible to centrally control all vehicles in the German road traffic system or all data packets in a large communication network like the Internet. Therefore, it is of interest to study selfish behavior in such shared networks. To do so, we use and extend two well-known classes of games that model selfish routing in networks: Wardrop games and congestion games.

Wardrop games were already studied in the 1920s by Pigou [77] and in the 1950s by Beckman et al. [12] and Wardrop [92] in the context of road traffic systems. In a Wardrop game, we are given a *network*, amounts of *traffic* between pairs of network nodes, and for each edge a *latency function* describing the time needed to traverse the edge depending on the total traffic on this edge. Traffic is modeled as *network flow*, i.e., the traffic associated with a pair of an origin and a destination node is allowed to split into arbitrary pieces. The cost that traffic on an origin-destination path experiences is given by the sum of edge latencies on this path. In a *Wardrop equilibrium*, no fraction of the traffic assigned to some path, however small, can decrease its cost by unilaterally switching to another path.

Originally introduced by Rosenthal [80] in the 1970s, *congestion games* can be used to model resource sharing among players. Since routing is a typical resource sharing problem, the restricted class of *network congestion games* is of interest. In such a game, the resources are the edges of a given *network* with *latency functions* on the edges. There is a *finite* number of *players* and each of them has an origin and a destination node in the network. Each player selects as its pure *strategy* one path from its origin to its destination node and experiences as its *private cost* the sum of edge latencies on this path. A selection of pure strategies is a pure *Nash equilibrium* [72, 73] if no player can decrease its private cost by unilaterally deviating to another path. Since Nash equilibria are stable states in which no player has an incentive to modify its strategy they are the most commonly used solution concept of rational selfish behavior in game theory.

In recent years, Wardrop and congestion games have attracted a great deal of attention as combining ideas from game theory and computer science has become increasingly popular. The attention in the theoretical computer science community was initiated by Koutsoupias and Papadimitriou [60] in 1999. They brought up the problem of bounding the so-called *price of anarchy* [76] that measures the worst-possible inefficiency of equilibria with respect to a social welfare measure. Later the *price of stability* [8] was introduced that measures the worst-possible inefficiency of best equilibria. Beyond the research on the degradation of social welfare due to selfish behavior, many recent papers (that we will discuss in Chapter 4) focus on other natural questions related to equilibria. These questions include the existence of equilibria, the complexity of computing equilibria, and the convergence to equilibria.

1.2 Outline of our Results

This thesis is primarily concerned with network routing games where the players have *incomplete information* since they do not know all relevant parameters of the game. As described earlier in the abstract, we use two different approaches to model incomplete information. On the one hand, we investigate network routing games where the players are uncertain about the edge latency functions and use, as in Milchtaich's [69] model, different *player-specific* latency functions for an edge. On the other hand, we use Harsanyi's [51] incomplete information concept to define so-called Bayesian routing games and study this scenario where the players do not know each other's weight.

Since we will give detailed overviews of our results in the Sections 5.1.1, 6.1.1, 7.1.1, 8.1.1, and 9.1.1 we only briefly sketch these results here:

- Congestion Games with Player-Specific Constants

  In Chapter 5, we consider congestion games with player-specific constants where players use the same *delay function* but different player-specific *constants* for each particular resource. The delay function and the player-specific constants constitute player-specific latency functions by means of an operation such as addition or multiplication.

We show that there are subclasses of these games where selfish steps of players can be used to reach a pure Nash equilibrium, whereas we prove for another subclass that a pure Nash equilibrium is guaranteed to exist although selfish step sequences can be of infinite length. Moreover, we show that under certain conditions it is PLS-hard [55] to compute a pure Nash equilibrium.

Although Milchtaich [69] earlier gave fundamental results on the convergence of selfish step sequences to pure Nash equilibria the results that we give in Chapters 5 and Chapter 6 allow to better understand under which conditions selfish steps can be used to reach an equilibrium.

- Congestion Games with Player-Specific Affine Latency Functions

  In Chapter 6, we study congestion games with player-specific *affine* latency functions.

  We characterize for which subclasses of these games selfish steps can be used to get a pure Nash equilibrium and we show that there is a game that does not have a pure Nash equilibrium. Furthermore, we prove an upper and an asymptotically tight lower bound on the price of anarchy.

  These and the results by Georgiou et al. [48] (see Section 4.3) are the first results on the price of anarchy for congestion games with player-specific latency functions.

- Wardrop Games with Player-Specific Affine Latency Functions

  In Chapter 7, we focus on a generalization of Wardrop games that allows for player-specific latency functions. Most of our results apply for the case where the player-specific latency functions are *affine*.

  For a specific subclass of these games, we prove that a Wardrop equilibrium can be computed in polynomial time by minimizing a new convex potential function that we introduce. However, we show that a similar argumentation cannot be applied for a certain more general setting. We also prove an upper and a lower bound on the price of anarchy.

  While Wardrop games have been studied extensively in recent years (see Section 4.4) these are the first results on the more general setting with player-specific latency functions.

- Bayesian Routing Games

  In Chapter 8, we consider Bayesian routing games with incomplete information where, in contrast to congestion games, the players do not know each other's weight. Following Harsanyi's approach [51], we use types and a probability distribution over all possible type profiles to model the players' uncertainty about each other's weight.

We show that every Bayesian routing game has a pure equilibrium. For a subclass of these games we give a polynomial-time algorithm to compute one. We also characterize equilibria that maximize the private costs of all players. This enables us to prove results on the price of anarchy for three different social welfare measures.

Although a lot of work has been done on congestion games (see Sections 4.1 and 4.2) we introduce and study with our Bayesian routing games the first model with incomplete information on the player weights.

- Bottleneck Games with Splittable Traffic

  In Chapter 9, we study, as a complete information setting, bottleneck games with splittable traffic. Here, the latency on a path is given by the *maximum* latency of an edge on this path whereas for Wardrop games it is given by the *sum* of these latencies.

  We characterize for which games the social welfare of Wardrop equilibria is unique. Moreover, we show that the price of stability is independent of the network topology and we give the exact price of stability for games with M/M/1 latency functions [57].

  These and the results of Cole et al. [22] (see Section 4.5) are the first results on a Wardrop-like maximum latency setting.

1.3 Publications

Results presented in this thesis are published in the Proceedings of the *International Colloquium on Automata, Languages, and Programming (ICALP)* [45, 47], the Proceedings of the *Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)* [46], the Proceedings of the *International Workshop on Internet and Network Economics (WINE)* [68], and will appear in the Proceedings of the *International Symposium on Mathematical Foundations of Computer Science* (MFCS) [66] and *Theory of Computing Systems (TOCS)* [46].

Since this thesis is on game theory it does not include the results on branch-and-bound algorithms for the test cover problem that I also published while I was working on my PhD project. These results appeared in the Proceedings of the *International Workshop on Efficient and Experimental Algorithms (WEA)* [32] and the *ACM Journal of Experimental Algorithmics (JEA)* [33]. Furthermore, this thesis does not include recent results on cost-sharing mechanisms that will appear in the Proceedings of the *International Symposium on Mathematical Foundations of Computer Science* (MFCS) [17].

Many results given in this thesis were developed in joint work with the co-authors of the aforementioned publications. Throughout this thesis I will give all proofs that I developed and also proofs that we obtained in collaboration.

Since the proofs to that I had no contribution are not included some theorems in this thesis are not followed by a proof.

1.4 Road Map

The rest of this thesis is organized as follows. In Chapter 2, we give some preliminaries, whereas we formally describe the models considered in this thesis in Chapter 3. We summarize previous work related to this thesis in Chapter 4. In the remaining chapters, we give our results on congestion games with player-specific constants (Chapter 5), congestion games with player-specific affine latency functions (Chapter 6), Wardrop games with player-specific affine latency functions (Chapter 7), Bayesian routing games (Chapter 8), and bottleneck games with splittable traffic (Chapter 9).

CHAPTER 2

Preliminaries

This chapter presents some basic notation (Section 2.1), introduces relevant graph types (Section 2.2), affine functions (Section 2.3), PLS-problems (Section 2.4), maximum flows (Section 2.5), and totally ordered abelian groups (Section 2.6).

2.1 Notation

Denote $[k] = \{1, \ldots, k\}$ for each integer $k \geq 0$. For a vector $\mathsf{v} = (v_1, \ldots, v_n)$ let $\mathsf{v_{-i}} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ and $(\mathsf{v_{-i}}, v_i') = (v_1, \ldots, v_{i-1}, v_i', v_{i+1}, \ldots, v_n)$.

2.2 Graphs and Networks

Throughout this thesis we will mainly consider games in networks with *directed edges* where there is for every player one *origin node* and one *destination node*. Such games can be interpreted as routing games. We will now introduce some important types of multigraphs / networks that will be used in this thesis.

2.2.1 Asymmetric and Symmetric Networks

If all players share the same origin node and the same destination node, then the network is *symmetric*, otherwise it is *asymmetric*. If a symmetric network $G = (V, E)$ is considered, we denote by $v_o \in V$ the origin node of all players and by $v_d \in V$ the destination node of all players.

2.2.2 Series Parallel Graphs

Series parallel graphs are sometimes also called two terminal series parallel graphs. Series parallel is a recursively defined property: As the *base case*, the graph that only consists of two nodes v_o, v_d and a single edge (v_o, v_d) is series parallel with terminals (v_o, v_d). An arbitrary multigraph G is series parallel with terminals (v_o, v_d) if it can be constructed from two series parallel graphs with terminals (v_o^1, v_d^1) and (v_o^2, v_d^2) connected either in series or in parallel. In a *series connection*, $v_d^1 = v_o^2$, $v_o = v_o^1$, and $v_d = v_d^2$. In a *parallel connection*, $v_o = v_o^1 = v_o^2$ and $v_d = v_d^1 = v_d^2$.

Note that all series parallel graphs are symmetric.

2.2.3 Parallel Link Graphs

A *parallel link graph* $G = (V, E)$ is a multigraph that only consists of the origin node $v_o \in V$ of all players, the destination node $v_d \in V$ of all players, and $m \in \mathbb{N}$ edges each of them connecting these two nodes. Obviously, all parallel link graphs are series parallel graphs.

2.3 Affine and Linear Functions

Throughout this thesis we call a function *affine* if it is of the form $f(x) = a \cdot x + b$, whereas we call a function *linear* if it is of the form $f(x) = a \cdot x$. Clearly, every linear function is affine. Although there is scientific literature that also calls functions linear that are of the form $a \cdot x + b$ where $b \neq 0$ we distinguish between affine and linear functions since this is advantageous for the presentation of our results. This distinction was also used by other authors (see e.g. [5, 23, 29]).

2.4 PLS(-complete) Problems

Johnson et al. [55] defined in the 1980s the complexity class *PLS* (Polynomial-time Local Search) that includes optimization problems where the goal is to find a *local optimum* for a given instance; this is a feasible solution with no feasible solution of better objective value in its well-determined neighborhood. A problem Π in PLS has an associated set of instances $\mathcal{I}_\Pi$. There is, for every instance $I \in \mathcal{I}_\Pi$, a set of feasible solutions $\mathcal{F}(I)$. Furthermore, there are three polynomial time algorithms A, B, and C.

- A computes for every instance I a feasible solution $S \in \mathcal{F}(I)$;

- B computes for a feasible solution $S \in \mathcal{F}(I)$ the objectice value;

- C determines, for a feasible solution $S \in \mathcal{F}(I)$, whether S is locally optimal and, if not, it outputs a feasible solution in the neighborhood of S with better objective value.

A PLS-problem Π_1 is *PLS-reducible* [55] to a PLS-problem Π_2 if there are two polynomial time computable functions F_1 and F_2 such that

- F_1 maps instances $I \in \mathcal{I}_{\Pi_1}$ to instances $F_1(I) \in \mathcal{I}_{\Pi_2}$ and
- F_2 maps every local optimum of the instance $F_1(I)$ to a local optimum of I.

A PLS-problem Π is *PLS-complete* [55] if every problem in PLS is PLS-reducible to Π.

2.5 Maximum Flows and Minimum Cuts

The maximum flow problem is a classical optimization problem with many applications (see e.g. [3]). An instance $(G, v_o, v_d, (k_e)_{e \in E})$ of the *maximum flow problem* consists of a directed graph $G = (V, E)$, a source node $v_o \in V$, a sink node $v_d \in V$, and a capacity $k_e \geq 0$ for each edge $e \in E$. A *feasible flow* $H = (h_e)_{e \in E}$ assigns to each edge $e \in E$ an edge flow $h_e \geq 0$ satisfying the capacity constraints, i.e., $h_e \leq k_e$. Furthermore, the flow conservation constraints are fulfilled at all nodes. We denote by $|H|$ the total amount of flow that is shipped from v_o to v_d. A *maximum flow* H is a feasible flow that maximizes the total amount of flow $|H|$ shipped from v_o to v_d.

The maximum-flow minimum-cut theorem states that the maximum amount of flow possible for an instance is equal to the capacity of a minimum v_o-v_d-cut. It is possible to represent such a cut by a partition of the nodes V into two subsets (S, T) where $v_o \in S$ and $v_d \in T$. In this work, however, we regard a *minimum cut* as a set of edges $D(S, T) \subseteq E$ where all edges $e \in D(S, T)$ leave the partition S, i.e., they are of the form (u, v) where $u \in S$ and $v \in T$.

Given a minimum cut $D(S, T)$ that belongs to a maximum flow H, all edges $e \in D(S, T)$ are fully saturated, i.e., $h_e = k_e$, and $\sum_{e \in D(S,T)} h_e = |H|$. Each path from v_o to v_d includes at least one edge $e \in D(S, T)$.

2.6 Totally Ordered Abelian Groups

A *group* $(G, \odot)$ consists of a *ground set* G together with a binary operation $\odot : G \times G \to G$; $\odot$ is *associative* and allows for an *identity element* and *inverses*. The group $(G, \odot)$ is *abelian* if $\odot$ is *commutative*. We will consider *totally ordered abelian groups* with a *total order* on G [49] which satisfies *translation invariance*: for all triples $r, s, t \in G$, if $r \leq s$ then $r \odot t \leq s \odot t$.

Examples of totally ordered, translation-invariant abelian groups include

(i) $(\mathbb{R}^+ \setminus \{0\}, \cdot)$ under the usual number-ordering,

(ii) $(\mathbb{R}, +)$ under the usual number-ordering, and

(iii) $(\mathbb{R}^2, +)$ under the lexicographic ordering on pairs of numbers.

CHAPTER 3

Models

We introduce in this chapter congestion games with player-specific latency functions (Section 3.1), Wardrop games with player-specific latency functions (Section 3.2), Bayesian routing games (Section 3.3), and bottleneck games with splittable traffic (Section 3.4).

3.1 Congestion Games with Player-Specific Latency Functions

Congestion games with player-specific latency functions were (for the special case of parallel links) originally introduced by Milchtaich [69]. As new subclasses of these games we will in this section introduce dominance congestion games and congestion games with player-specific constants. The even more restricted class of unweighted congestion games was originally introduced by Rosenthal [80].

We will investigate the game classes introduced here in Chapters 5 and 6.

3.1.1 Instances

A *weighted congestion game with player-specific latency functions* Γ is a tuple

$$\Gamma = \left(n, E, (w_i)_{i\in[n]}, (S_i)_{i\in[n]}, (f_{ie})_{i\in[n],e\in E}\right).$$

Here, n is the finite number of *players* and E is the finite set of *resources*. For every player $i \in [n]$, $w_i > 0$ is the *weight* and $S_i \subseteq 2^E$ is the *strategy set* of player i. Denote $S = S_1 \times \ldots \times S_n$. For every player i and resource $e \in E$, there is a non-decreasing *latency function* $f_{ie} : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ that player i assigns to e.

In an *unweighted* congestion game with player-specific latency functions, the weights of all players are equal. W.l.o.g. we assume in this case that $w_1 = \ldots = w_n = 1$. In a *network* congestion game with player-specific latency functions the strategy set S_i of a player i corresponds to all paths from his origin node to his destination node in a directed network $G = (V, E)$.

Fix a congestion game Γ with player-specific latency functions. Consider a pair of players $i \neq j$ and a pair of resources (e, e') where $e \neq e'$. Say that i *dominates* j for the ordered pair (e, e') if for every pair of positive numbers $x, y \in \mathbb{R}^+$, $f_{ie}(x) > f_{ie'}(y)$ implies $f_{je}(x) > f_{je'}(y)$. Intuitively, i dominates j for (e, e') if the preference of i to switch strategy from e to e' always implies a corresponding preference for j. A congestion game with player-specific latency functions is called *dominance congestion game* if for all pairs of players $i \neq j$ and for all pairs of resources (e, e') where $e \neq e'$, either i dominates j for (e, e') or j dominates i for (e, e').

Fix a totally ordered, translation-invariant abelian group $(G, \odot)$, $G \subseteq \mathbb{R}$, under the usual number-ordering. A *congestion game with player-specific constants* is a congestion game Γ with player-specific latency functions such that

(i) for each resource $e \in E$, there is a non-decreasing *delay function* $g_e : \mathbb{R}_0^+ \to \mathbb{R}_0^+$, and

(ii) for each pair of a player $i \in [n]$ and a resource $e \in E$, there is a player-specific constant $c_{ie} > 0$,

so that for each player $i \in [n]$ and resource $e \in E$, $f_{ie}(x) = c_{ie} \odot g_e(x)$. In a *congestion game with player-specific additive constants* (resp., *player-specific multiplicative constants*), G is $\mathbb{R}$ and $\odot$ is $+$ (resp., G is $\mathbb{R}^+ \setminus \{0\}$ and $\odot$ is $\cdot$).

If all latency functions f_{ie} are of the form $f_{ie}(x) = g_e(x)$ then Γ is a *congestion game* and there is only one latency function $g_e : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ for an edge $e \in E$ that is used by all players.

3.1.2 Affine Latency Functions

We will sometimes consider congestion games with player-specific *affine* latency functions $f_{ie}(x) = a_{ie} \cdot x + b_{ie}$ where $a_{ie}, b_{ie} \geq 0$. For such a game Γ with affine latency functions denote

$$\Delta(\Gamma) \quad = \quad \max_{e \in E;\, i,k \in [n]} \left\{ \frac{a_{i\,e}}{a_{k\,e}} ; a_{i\,e} < \infty, a_{k\,e} < \infty \right\}$$

with the understanding that $\frac{0}{0} = 1$. $\Delta(\Gamma)$ describes the maximum factor by which the slopes of the player-specific affine latency functions for one edge differ. Note that $\Delta(\Gamma)$ does not depend on the constants b_{ie} of the latency functions. We will also sometimes consider congestion games with player-specific *linear* latency functions where all latency functions are of the form $f_{i\,e}(x) = a_{i\,e} \cdot x, a_{i\,e} > 0$.

3.1.3 Strategies and Strategy Profiles

A *pure strategy* for player $i \in [n]$ is some specific $s_i \in S_i$ whereas a *mixed strategy* $Q_i = (q(i, s_i))_{s_i \in S_i}$ is a probability distribution over S_i, where $q(i, s_i)$ denotes the probability that player i chooses the pure strategy s_i. A *pure strategy profile* is an n-tuple $\mathsf{s} = (s_1, \ldots, s_n) \in S$ whereas a *mixed strategy profile* $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is represented by an n-tuple of mixed strategies. Clearly, every pure strategy (profile) is a mixed strategy (profile). For a mixed strategy profile $\mathbf{Q}$ and a pure strategy profile s define $q(\mathsf{s})$ as the probability that s is selected, i.e.,

$$q(\mathsf{s}) \quad = \quad \prod_{i \in [n]} q(i, s_i).$$

3.1.4 Private Cost

Fix any pure strategy profile s, and denote the *load* on resource $e \in E$ by

$$\delta_e(\mathsf{s}) \quad = \quad \sum_{i \in [n], s_i \ni e} w_i.$$

The *private cost* of player $i \in [n]$ is defined by

$$\mathsf{PC}_i(\mathsf{s}) \quad = \quad \sum_{e \in s_i} f_{ie}\left(\delta_e(\mathsf{s})\right).$$

For a mixed strategy profile $\mathbf{Q}$, the *private cost* of player $i \in [n]$ is

$$\mathsf{PC}_i(\mathbf{Q}) \quad = \quad \sum_{\mathsf{s} \in S} q(\mathsf{s}) \cdot \mathsf{PC}_i(\mathsf{s}).$$

3.1.5 Nash Equilibria and Selfish Steps

We are interested in a special class of strategy profiles called Nash equilibria that we describe here. Given a game and an associated strategy profile, a player $i \in [n]$ is *satisfied* if it can not improve its private cost by unilaterally changing its strategy. Otherwise, player i is *unsatisfied*. The mixed strategy profile $\mathbf{Q}$ is a *Nash equilibrium* if and only if all players $i \in [n]$ are satisfied, i.e., $\mathsf{PC}_i(\mathbf{Q}) \leq \mathsf{PC}_i(\mathbf{Q}_{-i}, s_i)$ for all $i \in [n]$ and all $s_i \in S_i$. Depending on the type of the strategy profile $\mathbf{Q}$, we distinguish between *pure* and *mixed* Nash equilibria.

Fix any pure strategy profile. In a *selfish step*, exactly one unsatisfied player is allowed to change its pure strategy such that its private cost decreases. A selfish step is *greedy* if the changing player is satisfied after he has done the selfish step.

A game Γ possesses the *finite best-reply property* if any sequence of greedy selfish steps is finite. If even any sequence of selfish steps is finite it possesses in addition the *finite improvement property.*

A *potential function*[1] for the game Γ is a function $\Phi : S_1 \times \ldots \times S_n \to \mathbb{R}$ that decreases when a player takes a selfish step. If a game has a potential function then it also has the finite improvement property. Note, that the finite improvement property implies the finite best-reply property which again implies the existence of a pure Nash equilibrium [70].

3.1.6 Social Cost and Price of Anarchy

Associated with a congestion game Γ and a mixed strategy profile $\mathbf{Q}$ is the *social cost* as a measure of social welfare. We consider social cost as the expected total latency [85], i.e.,

$$\begin{aligned} \mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma) &= \sum_{i\in[n]} w_i \cdot \mathsf{PC}_i(\mathbf{Q}) \\ &= \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{i\in[n]} w_i \cdot \sum_{e\in s_i} f_{ie}\left(\delta_e(\mathsf{s})\right). \end{aligned}$$

The *optimum social cost* associated with a game Γ is defined by

$$\mathsf{OPT}_{\mathsf{TL}}(\Gamma) = \min_{\mathsf{s}\in S} \mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma).$$

The *price of anarchy* denoted by $\mathsf{PoA}_{\mathsf{TL}}$ measures how different the optimum social cost and the social cost of mixed Nash equilibria can be. It is given by

$$\mathsf{PoA}_{\mathsf{TL}} = \sup_{\Gamma,\,\mathbf{Q}\text{ is NE}} \frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)}.$$

3.2 Wardrop Games with Player-Specific Latency Functions

In this section, we introduce the new class of Wardrop games with player-specific latency functions. In contrast to traditional Wardrop games [12, 77, 92] these games allow latency functions that are player-specific.

We will study Wardrop games with player-specific latency functions in Chapter 7.

[1] Such a potential function is called generalized ordinal potential in the paper of Monderer and Shapley [70].

3.2.1 Instances

A *Wardrop game with player-specific latency functions* Υ is a tuple

$$\Upsilon = (n, G, v_o, v_d, (w_i)_{i\in[n]}, (f_{ie})_{i\in[n],e\in E})$$

Here, n is the number of *players* and $G = (V, E)$ is a directed symmetric *multigraph* with an origin $v_o \in V$ and a destination $v_d \in V$. For every player $i \in [n]$, w_i is the *traffic* of player i. Edge latency functions $f_{i\,e}$ are player-specific and $f_{i\,e} : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ is the non-negative, non-decreasing, and continuous *player-specific latency function* that player $i \in [n]$ assigns to edge $e \in E$.

3.2.2 Affine Latency Functions

In the majority of cases, we consider player-specific *affine* latency functions $f_{ie}(u) = a_{ie} \cdot u + b_{ie}$ with $a_{ie}, b_{ie} \geq 0$. For a Wardrop game Υ with affine latency functions denote

$$\Delta(\Upsilon) \;=\; \max_{e\in E;\, i,k\in[n]} \left\{ \frac{a_{i\,e}}{a_{k\,e}}; a_{i\,e} < \infty, a_{k\,e} < \infty \right\}$$

with the understanding that $\frac{0}{0} = 1$. We will also study Wardrop games with player-specific *linear* latency functions where all latency functions are of the form $f_{i\,e}(u) = a_{i\,e} \cdot u, a_{i\,e} > 0$.

3.2.3 Strategies and Strategy Profiles

Let $\mathcal{P}$ be the set of all paths from the origin node v_o to the destination node v_d. A player $i \in [n]$ can split its traffic w_i over the paths in $\mathcal{P}$. A (pure) *strategy* for player $i \in [n]$ is a tuple $\mathsf{x}_i = (x_{iP})_{P\in\mathcal{P}}$ with $\sum_{P\in\mathcal{P}} x_{iP} = w_i$ and $x_{iP} \geq 0$ for all $P \in \mathcal{P}$. A *strategy profile* $\mathbf{x} = (\mathsf{x}_1, \ldots, \mathsf{x}_n)$ is an n-tuple of strategies for the players.

3.2.4 Wardrop Equilibria

For a strategy profile $\mathbf{x}$ the load $\delta_e(\mathbf{x})$ on an edge $e \in E$ is given by

$$\delta_e(\mathbf{x}) \;=\; \sum_{i\in[n]} \sum_{P\in\mathcal{P}, P\ni e} x_{iP}.$$

A strategy profile $\mathbf{x}$ is a *Wardrop equilibrium*, if for every player $i \in [n]$ and every $P, P' \in \mathcal{P}$ with $x_{iP} > 0$ it holds that

$$\sum_{e\in P} f_{i\,e}(\delta_e(\mathbf{x})) \;\leq\; \sum_{e\in P'} f_{i\,e}(\delta_e(\mathbf{x})).$$

Observe that in a Wardrop equilibrium all flow paths of a player have equal latency. We can regard each player $i \in [n]$ as a service provider who has many clients each handling a negligible small amount of traffic. In a Wardrop equilibrium each service provider satisfies all his clients because none of them can improve its experienced latency.

3.2.5 Social Cost and Price of Anarchy

If a Wardrop game Υ with player-specific latency functions and a strategy profile $\mathbf{x}$ are given the *social cost* $\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon)$ is defined by:

$$\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon) \quad = \quad \sum_{i \in [n]} \sum_{P \in \mathcal{P}} x_{iP} \cdot \sum_{e \in P} f_{i\,e}(\delta_e(\mathbf{x})).$$

This social cost measure is motivated by the interpretation of Υ as a game with infinitely many (sub-)players with negligible demand and models the sum of the players latencies. The *optimum social cost* $\mathsf{OPT}_{\mathsf{TL}}(\Upsilon)$ associated with a Wardrop game Υ with player-specific latency functions is given by the smallest possible social cost of a strategy profile, i.e.,

$$\mathsf{OPT}_{\mathsf{TL}}(\Upsilon) = \min_{\mathbf{x}} \mathsf{SC}_{\mathsf{TL}}(\mathbf{x}).$$

The *price of anarchy* denoted by $\mathsf{PoA}_{\mathsf{TL}}$ measures how different the optimum social cost and the social cost of Wardrop equilibria can be. It is given by

$$\mathsf{PoA}_{\mathsf{TL}} \quad = \quad \sup_{\Upsilon,\ \mathbf{x} \text{ is WE}} \frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon)}{\mathsf{OPT}_{\mathsf{TL}}(\Upsilon)}.$$

3.3 Bayesian Routing Games

In this section, we introduce the new class of Bayesian routing games. These games are based on Harsanyi's [51] concept of Bayesian games.

We will consider Bayesian routing games in Chapter 8.

3.3.1 Instances

A *Bayesian routing game* is a tuple $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$. Each of n players $1, 2, \ldots, n$ wishes to assign a particular amount of weight to one of m *links* $1, 2, \ldots, m$. Throughout, we assume that $n \geq 2$ and $m \geq 2$. Denote $\mathbf{c} = (c_1, \ldots, c_m)$, where $c_j > 0$ is the *capacity* of link $j \in [m]$. In the case of *identical links*, all capacities equal 1. In this case, we write $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$. Link capacities vary arbitrarily in the case of *related links*. For each player $i \in [n]$, there is a finite set of possible types T_i; for each type $t \in T_i$, denote by $w(t)$ the *weight* of type t, $w(t) \geq 0$. Denote $T = T_1 \times \ldots \times T_n$, the set of all possible *type profiles*. For simplicity we assume that the weights $(w(t_i))_{t_i \in T_i, i \in [n]}$ are encoded in T, so we do not include them in the game tuple. We use the term *type agent* (i, t) to refer to the type $t \in T_i$ of player $i \in [n]$.

There is a joint probability distribution $\mathbf{p} = (p(t_1, \ldots, t_n))_{(t_1, \ldots, t_n) \in T}$, called *type distribution*, over the set of type profiles T; thus, $\mathbf{p}$ is a function $\mathbf{p} : T \to [0, 1]$ and $\sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) = 1$. Denote by $p(i, t)$ the probability that player i is of type t; so,

$$p(i, t) \;=\; \sum_{(t_1, \ldots, t_n) \in T : t_i = t} p(t_1, \ldots, t_n).$$

We say that $\mathbf{p}$ is *independent* if

$$p(t_1, \ldots, t_n) \;=\; \prod_{i \in [n]} p(i, t_i) \qquad \text{for all } (t_1, \ldots, t_n) \in T,$$

otherwise, $\mathbf{p}$ is *correlated*. By the definition of conditional probability,

$$p(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n | t_k = t) \;=\; \frac{p(t_1, \ldots, t_{k-1}, t, t_{k+1}, \ldots, t_n)}{p(k, t)};$$

that is, the probability of a type profile $(t_1, \ldots, t_n)$ given that $t_k = t$ is the probability of type profile $(t_1, \ldots, t_n)$ divided by the probability that player k is of type t. We only consider instances where $p(k, t) > 0$ for all players $k \in [n]$ and all types $t \in T_k$. Denote by $W(i)$ the *expected weight* of player $i \in [n]$; clearly,

$$\begin{aligned} W(i) &= \sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) \cdot w(t_i) \\ &= \sum_{t \in T_i} p(i, t) \cdot w(t). \end{aligned}$$

Furthermore, define the *expected total weight* as

$$W \;=\; \sum_{i \in [n]} W(i).$$

For any pair of players $i, s \in [n]$ and for any type $t \in T_i$, define $W(s | t_i = t)$ as the *conditional expected weight* of player s, given that player i has type t; so,

$$W(s | t_i = t) \;=\; \sum_{\substack{(t_1, \ldots, t_n) \in T: \\ t_i = t}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = t) \cdot w(t_s).$$

For the case of independent type distribution we have $W(s | t_i = t) = W(s)$ for all types $t \in T_i$ of player i.

A special instance of our Bayesian routing game in which each player has only a single type is a *complete information routing game* or *weighted congestion game on parallel links* (see Section 3.1.1). For such a game, we write $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T, 1)$. Here, the set T contains only one type profile t that is used with probability 1.

3.3.2 Strategies and Strategy Profiles

A *pure strategy* σ_i for player $i \in [n]$ is a mapping of the set of possible types T_i to the set of links $[m]$; so σ_i is a function $\sigma_i : T_i \rightarrow [m]$. Denote as Σ_i the set of all possible pure strategies for player $i \in [n]$; denote $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$. A *mixed strategy* $Q_i = (q(i, \sigma_i))_{\sigma_i \in \Sigma_i}$ for player $i \in [n]$ is a probability distribution over Σ_i; here, $q(i, \sigma_i)$ denotes the probability that player i chooses the pure strategy σ_i.

The *support* of a mixed strategy Q_i for player $i \in [n]$, denoted $\mathsf{support}_{Q_i}(i)$, is the set of links to which player i assigns at least one type $t \in T_i$ with positive probability, that is,

$$\mathsf{support}_{Q_i}(i) \quad = \quad \{j \in [m] \mid \exists \sigma_i \in \Sigma_i, \exists t \in T_i \text{ with } q(i, \sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

Similarly, the support of any type $t \in T_i$ of player $i \in [n]$ is defined by

$$\mathsf{support}_{Q_i}(t) \quad = \quad \{j \in [m] \mid \exists \sigma_i \in \Sigma_i \text{ with } q(i, \sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

Note that

$$\mathsf{support}_{Q_i}(i) \quad = \quad \bigcup_{t \in T_i} \mathsf{support}_{Q_i}(t).$$

A *pure strategy profile* σ is an n-tuple $(\sigma_1, \ldots, \sigma_n) \in \Sigma$. Call σ *normal* if $\sigma_i(t) = \sigma_i(t')$ for all types $t, t' \in T_i$ and for all players $i \in [n]$. So, each player $i \in [n]$ does not distinguish among its types in a normal pure strategy profile.

A *mixed strategy profile* $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is an n-tuple of mixed strategies. Call a mixed strategy profile $\mathbf{F} = (F_1, \ldots, F_n)$ *fully mixed* if each player assigns strictly positive probability to each of its pure strategies; that is, $f(i, \sigma_i) > 0$ for all players $i \in [n]$ and all strategies $\sigma_i \in \Sigma_i$. Notice that $\mathsf{support}_{F_i}(i) = [m]$ for all players $i \in [n]$ and $\mathsf{support}_{F_i}(t) = [m]$ for all players $i \in [n]$ and types $t \in T_i$.

3.3.3 Private Cost

Pure Strategy Profiles

Fix any type distribution $\mathbf{p}$ and a pure strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$. The *expected load* on link $j \in [m]$, denoted $\delta_j(\sigma, \mathbf{p})$, is defined by

$$\delta_j(\sigma, \mathbf{p}) \quad = \quad \sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) \cdot \sum_{\substack{i \in [n]: \\ \sigma_i(t_i) = j}} w(t_i).$$

In the same way, denote as $\delta_j^{-k}(\sigma, (\mathbf{p}|t_k = t))$ the *conditional expected load* of all players $i \in [n]$ other than k on link $j \in [m]$ given that $t_k = t$; so,

$$\delta_j^{-k}(\sigma, (\mathbf{p}|t_k = t)) \quad = \quad \sum_{\substack{(t_1, \ldots, t_n) \in T: \\ t_k = t}} p(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n | t_k = t) \sum_{\substack{i \in [n] \setminus \{k\}: \\ \sigma_i(t_i) = j}} w(t_i).$$

Let $\lambda^j_{(i,t)}(\sigma, \mathbf{p})$ be the private cost of type agent (i,t) when it is assigned to link $j \in [m]$; so,

$$\lambda^j_{(i,t)}(\sigma, \mathbf{p}) \quad = \quad \frac{\delta_j^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)}{c_j}.$$

Denote as $v_{(i,t)}(\sigma, \mathbf{p})$ the *conditional private cost* of player $i \in [n]$, given that player i is of type t; this is also the private cost of type agent (i,t); so

$$v_{(i,t)}(\sigma, \mathbf{p}) \quad = \quad \lambda^{\sigma_i(t)}_{(i,t)}(\sigma, \mathbf{p}).$$

Note that $v_{(i,t)}(\sigma, \mathbf{p})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player i. Finally, denote as $\mathsf{PC}_i(\sigma, \mathbf{p})$ the *private cost* of player i; clearly,

$$\mathsf{PC}_i(\sigma, \mathbf{p}) \quad = \quad \sum_{t \in T_i} p(i,t) \cdot v_{(i,t)}(\sigma, \mathbf{p}).$$

Mixed Strategy Profiles

Fix any type distribution $\mathbf{p}$ and a mixed strategy profile $\mathbf{Q}$. The *expected load* on link $j \in [m]$, denoted $\delta_j(\mathbf{Q}, \mathbf{p})$, is defined by

$$\delta_j(\mathbf{Q}, \mathbf{p}) \quad = \quad \sum_{\sigma \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \cdot \delta_j(\sigma, \mathbf{p}).$$

In the same way, denote as $\delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t))$ the *conditional expected load* of all players $i \in [n]$ other than k on link $j \in [m]$ given that $t_k = t$; so,

$$\delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t)) \quad = \quad \sum_{\sigma \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \cdot \delta_j^{-k}(\sigma, (\mathbf{p}|t_k = t)).$$

For the case of an independent type distribution $\mathbf{p}$, we get that for all types $t, t' \in T_k$, $\delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t)) = \delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t'))$. Therefore, to simplify notation, we write in this case $\delta_j^{-k}(\mathbf{Q}, \mathbf{p})$.

Let $\lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p})$ be the private cost of type agent (i,t) when it is assigned to link $j \in [m]$; so,

$$\lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p}) \quad = \quad \frac{\delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) + w(t)}{c_j}.$$

Denote as $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ the *conditional private cost* of player $i \in [n]$, given that player i is of type t; this is also the private cost of type agent (i,t); so,

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \quad = \quad \sum_{\sigma_i \in \Sigma_i} q(i, \sigma_i) \cdot \lambda^{\sigma_i(t)}_{(i,t)}(\mathbf{Q}, \mathbf{p}).$$

Note that $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player i. Finally, denote as $\mathsf{PC}_i(\mathbf{Q}, \mathbf{p})$ the *private cost* of player i; clearly,

$$\mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) \quad = \quad \sum_{t \in T_i} p(i,t) \cdot v_{(i,t)}(\mathbf{Q}, \mathbf{p}).$$

3.3.4 Bayesian Nash Equilibria

A strategy profile is a Bayesian Nash equilibrium, if no player has an incentive to deviate from its strategy; that is, no player can possibly decrease its private cost when other players are sticking to their strategies. Formally, the mixed strategy profile $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is a *Bayesian Nash equilibrium* if

$$\mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) \ \leq \ \mathsf{PC}_i(\mathbf{Q}', \mathbf{p})$$

for all mixed strategy profiles $\mathbf{Q}' = (Q_1, \ldots, Q_{i-1}, Q'_i, Q_{i+1}, \ldots, Q_n)$ and for all players $i \in [n]$. Moreover, since $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player i, the above condition is equivalent to

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \ \leq \ v_{(i,t)}(\mathbf{Q}', \mathbf{p})$$

for all mixed strategy profiles $\mathbf{Q}' = (Q_1, \ldots, Q_{i-1}, Q'_i, Q_{i+1}, \ldots, Q_n)$ and for all players $i \in [n]$ and types $t \in T_i$. Note that $\mathbf{Q}$ is a Bayesian Nash equilibrium if and only if for all players $i \in [n]$ and types $t \in T_i$,

$$\begin{aligned} v_{(i,t)}(\mathbf{Q}, \mathbf{p}) &= \lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p}), \quad \text{for } j \in \mathsf{support}_{Q_i}(t), \text{ and} \\ v_{(i,t)}(\mathbf{Q}, \mathbf{p}) &\leq \lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p}), \quad \text{for } j \notin \mathsf{support}_{Q_i}(t). \end{aligned}$$

We refer to these conditions as the *Bayesian Nash equilibrium conditions.*

3.3.5 Social Cost and the Price of Anarchy

Associated with a Bayesian routing game $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ and a mixed strategy profile $\mathbf{Q}$ is the *social cost* as a measure of social welfare. We consider three different measures for social cost.

The *expected maximum latency*, which is the expectation over all player choices and type profiles, of the maximum latency on a link; so

$$\begin{aligned} &\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma) \\ &= \sum_{(\sigma_1,\ldots,\sigma_n)\in\Sigma} \prod_{i\in[n]} q(i,\sigma_i) \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i) \right\} \\ &= \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \sum_{(\sigma_1,\ldots,\sigma_n)\in\Sigma} \prod_{i\in[n]} q(i,\sigma_i) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i) \right\}. \end{aligned}$$

The *sum of private costs* is given by

$$\mathsf{SC}_{\mathsf{SUM}}(\mathbf{Q}, \Gamma) \ = \ \sum_{i\in[n]} \mathsf{PC}_i(\mathbf{Q}, \mathbf{p}).$$

And the *maximum private cost* is defined by

$$\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma) \quad = \quad \max_{i\in[n]} \mathsf{PC}_i(\mathbf{Q},\mathbf{p}).$$

Let $* \in \{\mathsf{MSP},\mathsf{SUM},\mathsf{MAX}\}$. Denote the corresponding *optimum social cost* by $\mathsf{OPT}_*(\Gamma) = \min_{\mathbf{Q}} \mathsf{SC}_*(\mathbf{Q},\Gamma)$. The *price of anarchy* PoA_* is the supremum, over all instances Γ and Bayesian Nash equilibria $\mathbf{Q}$, of the ratio $\frac{\mathsf{SC}_*(\mathbf{Q},\Gamma)}{\mathsf{OPT}_*(\Gamma)}$; that is,

$$\mathsf{PoA}_* \quad = \quad \sup_{\Gamma,\,\mathrm{BNE}\,\mathbf{Q}} \frac{\mathsf{SC}_*(\mathbf{Q},\Gamma)}{\mathsf{OPT}_*(\Gamma)}.$$

3.3.6 Weighted Bayesian Congestion Games

A generalization of the Bayesian routing game considered in our work is the *weighted Bayesian congestion game with affine latency functions*. Like in a classical congestion game (see Section 3.1), each player $i \in [n]$ can be assigned to a subset s_i of the resources out of a given set $S_i \subseteq 2^{[m]}$ of subsets of resources. The latency function of resource $e \in [m]$ is given by an arbitrary, non-decreasing affine function $g_e(x) = a_e x + b_e$. For a Bayesian congestion game, a pure strategy profile σ is defined by $\sigma = (\sigma_1, \ldots, \sigma_n)$ with $\sigma_i : T_i \to S_i$ for all $i \in [n]$. Thus a pure strategy of a player can consist of *many* resources whereas in a Bayesian routing game a pure strategy is *one* link.

For a pure strategy profile σ, the *conditional expected load* of all players $i \in [n]$ other than k, on resource $e \in [m]$ given that $t_k = t$ is then

$$\delta_e^{-k}(\sigma, (\mathbf{p}|t_k = t)) \quad = \quad \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_k=t}} p(t_1,\ldots,t_{k-1},t_{k+1},\ldots,t_n|t_k=t) \sum_{\substack{i\in[n]\setminus\{k\}:\\ e\in\sigma_i(t_i)}} w(t_i);$$

whereas the conditional private cost of player i, given that player i is of type $t \in T_i$ is then defined by

$$v_{(i,t)}(\sigma,\mathbf{p}) \quad = \quad \sum_{e\in\sigma_i(t)} g_e(\delta_e^{-i}(\sigma,(\mathbf{p}|t_i=t)) + w(t)).$$

3.4 Bottleneck Games with Splittable Traffic

In this section, we introduce the new class of bottleneck games with splittable traffic. Here, the latency on a path is given by the *maximum* latency of an edge on this path whereas for the well-known Wardrop games [12, 77, 92] it is given by the *sum* of these latencies. Recently, Cole et al. [22] defined and studied a model that is very similar to our bottleneck games with splittable traffic (see Section 4.5 for a comparison).

We will investigate bottleneck games with splittable traffic in Chapter 9.

3.4.1 Instances

A *bottleneck game with splittable traffic* is a tuple

$$\Gamma \quad = \quad (G, v_o, v_d, (f_e)_{e \in E}, r)$$

where a *traffic* of $r \in \mathbb{R}^+$ has to be routed from an origin node v_o to a destination node v_d in a *network* $(G, s, t, (f_e)_{e \in E})$. Here, $G = (V, E)$ is a directed symmetric *multigraph*, $v_o, v_d \in V$ are distinct *origin* and *destination* nodes, and the f_e's are latency functions. Each of these functions $f_e : \mathbb{R}_0^+ \to \mathbb{R}_0^+ \cup \{\infty\}$ is non-negative, continuous, and non-decreasing.

For a non-empty set $\mathcal{F}$ of latency functions we define $\mathscr{G}(\mathcal{F})$ as the set of all bottleneck games with latency functions drawn from $\mathcal{F}$. The subset $\mathscr{P}(\mathcal{F}) \subset \mathscr{G}(\mathcal{F})$ consists of all games in $\mathscr{G}(\mathcal{F})$ that are defined on a graph of parallel links. To further differentiate we denote by $\mathscr{G}(\mathcal{F}, m, r) \subset \mathscr{G}(\mathcal{F})$ the set of games with at most m edges and a traffic of at most r. Likewise, $\mathscr{P}(\mathcal{F}, m, r) = \mathscr{G}(\mathcal{F}, m, r) \cap \mathscr{P}(\mathcal{F})$.

3.4.2 M/M/1 latency function

For ease of notation, we write $\Gamma = (G, v_o, v_d, (c_e)_{e \in E}, r)$ for a bottleneck game with splittable traffic and *M/M/1 latency functions* where $c_e > 0$ is the *capacity* for edge $e \in E$. The M/M/1 latency functions f_e, $e \in E$, are implicitly defined by

$$f_e(u) = \begin{cases} \frac{1}{c_e - u} & \text{if } u < c_e \\ \infty & \text{otherwise.} \end{cases}$$

Observe that the latency $f_e(u)$ approaches ∞ as the load u approaches c_e. We denote by $\mathcal{M}$ the set of all M/M/1 latency functions and by $\mathcal{M}_{\geq c} \subset \mathcal{M}$ the functions with a capacity of at least c where $c > 0$.

3.4.3 Strategy Profiles

The traffic r can split arbitrarily over the set $\mathcal{P}_{v_o v_d}$ of all possible *simple* paths from the origin v_o to the destination v_d. A *strategy profile* is a vector $\mathbf{x} = (x_P)_{P \in \mathcal{P}_{v_o v_d}}$ where $\sum_{P \in \mathcal{P}_{v_o v_d}} x_P = r$ and $x_P \geq 0$ for all $P \in \mathcal{P}_{v_o v_d}$. The *load* δ_e on an edge $e \in E$ is given by

$$\delta_e(\mathbf{x}) \quad = \quad \sum_{P \in \mathcal{P}_{v_o v_d}, P \ni e} x_P.$$

3.4.4 Wardrop Equilibria

A strategy profile is a *Wardrop equilibrium* if the bottleneck latency of each used path is not larger than the bottleneck latency of any other path, i.e., if for all $P, R \in \mathcal{P}_{v_o v_d}$

$$x_P > 0 \quad \Rightarrow \quad \max_{e \in P} f_e(\delta_e(\mathbf{x})) \leq \max_{e \in R} f_e(\delta_e(\mathbf{x})).$$

3.4.5 Social Cost and Price of Anarchy

The *social cost* of a strategy profile $\mathbf{x}$ is defined as the "canonically" weighted sum of all path latencies, i.e.,

$$\mathsf{SC}(\mathbf{x}, \Gamma) = \sum_{P \in \mathcal{P}_{v_o v_d}} x_P \cdot \max_{e \in P} f_e(\delta_e(\mathbf{x})).$$

If $\mathbf{x}$ is a Wardrop equilibrium, $l(\mathbf{x}) = \frac{\mathsf{SC}(\mathbf{x},\Gamma)}{r}$ denotes the unique bottleneck latency of all paths with non-zero flow. The *optimum* associated with a bottleneck game with splittable traffic Γ is the minimum social cost of any strategy profile: $\mathsf{OPT}(\Gamma) = \min_{\mathbf{x}} \mathsf{SC}(\mathbf{x}, \Gamma)$. The *price of anarchy* (PoA) and *price of stability* (PoS) for a set $\mathscr{G}$ of games are defined as

$$\mathsf{PoA}(\mathscr{G}) = \sup_{\substack{\Gamma \in \mathscr{G}, \\ \mathbf{x} \text{ Wardr. Equ. in } \Gamma}} \frac{\mathsf{SC}(\mathbf{x}, \Gamma)}{\mathsf{OPT}(\Gamma)},$$

$$\mathsf{PoS}(\mathscr{G}) = \sup_{\Gamma \in \mathscr{G}} \inf_{\mathbf{x} \text{ Wardr. Equ. in } \Gamma} \frac{\mathsf{SC}(\mathbf{x}, \Gamma)}{\mathsf{OPT}(\Gamma)},$$

where by definition $\frac{\infty}{\infty} = 1$ and $\frac{0}{0} = 1$. Furthermore, $\frac{u}{0} = \infty$ if $u > 0$.

3.4.6 Capacity of a Network

For a given network $(G, v_o, v_d, (f_e)_{e \in E})$ its *capacity* is given by

$$\begin{aligned} & C(G, v_o, v_d, (f_e)_{e \in E}) \\ & = \sup \left\{ r \in \mathbb{R}_0^+ \;\middle|\; \begin{array}{l} \exists \text{ strategy profile } \mathbf{x} \text{ with } \mathsf{SC}(\mathbf{x}, \Gamma) < \infty \\ \text{for } (G, v_o, v_d, (f_e)_{e \in E}, r) \end{array} \right\} \cup \{0\}. \end{aligned}$$

Whenever a bottleneck game with splittable traffic and M/M/1 latency functions on m parallel links is considered we assume that $c_1 \geq \ldots \geq c_m$ and denote

$$C = \sum_{i=1}^{m} c_i, \qquad C^{\leq i} = \sum_{k=1}^{i} c_k.$$

Clearly, C is just the capacity of the network.

CHAPTER 4

Related Work

In this chapter, we cite previous work that is directly related to this thesis. For a general introduction to game theory see [65, 71, 75].

4.1 Congestion games

Unweighted congestion games were introduced by Rosenthal [80, 81] and extensively studied afterwards (see [45, 90] for recent surveys). We gave a formal definition of these games in Section 3.1.

Existence and Computation of Equilibria

Rosenthal used a potential function to show that every unweighted congestion game possesses a pure Nash equilibrium [80]. Subsequent papers [70, 91] characterized games that admit a potential function as potential games and showed their relation to congestion games. Fotakis et al. [40] proved that every weighted asymmetric network congestion game with affine latency functions has a potential function and a pure Nash equilibrium; in contrast, there are weighted symmetric network congestion games with non-affine latency functions that have no pure Nash equilibrium even if there are only 2 players [40, 62]. Nevertheless, the existence of pure Nash equilibria is guaranteed whenever a weighted congestion game is considered where the strategy set of each player consists of the bases of a matroid on the set of resources [2]. It is strongly NP-complete to determine whether a given weighted network congestion game has a pure Nash equilibrium [28].

A polynomial time computation of a pure Nash equilibrium for an unweighted symmetric network congestion game is possible by reduction to the min-cost flow problem [30]. However, the problem becomes PLS-complete for both (non-network) unweighted symmetric congestion games [30] and unweighted asymmetric network congestion games where the edges of the network are either directed [30] or undirected [1].

Price of Anarchy

The price of anarchy was studied for congestion games with social cost defined as the total latency. For affine latency functions and both pure and mixed equilibria, it is exactly $\frac{5}{2}$ for unweighted [21] and $\frac{1}{2} \cdot (3 + \sqrt{5})$ for weighted congestion games [9]. The exact price of anarchy is also known for polynomial latency functions with non-negative coefficients [4].

4.2 Congestion Games on Parallel Links

The special case of congestion games on parallel links has attracted a lot of attention in the last couple of years; see, for example, [26, 34, 35, 38, 44, 59, 60, 64].

Existence and Computation of Equilibria

Weighted congestion games on parallel links have the finite improvement property (and hence a pure Nash equilibrium) if all latency functions are non-decreasing; in this setting, [38] implies that a pure Nash equilibrium can be computed in polynomial time by using the classical LPT algorithm due to Graham [50].

Fully Mixed Nash Equilibria

The fully mixed Nash equilibrium conjecture states for weighted congestion games on parallel links that the fully mixed Nash equilibrium has worst social cost among all Nash equilibria; it was motivated by some results in [67], explicitly formulated in [44], and further studied in [64]. For social cost defined as the sum of private costs, the conjecture holds [43, 63]; it was recently disproved for social cost defined as the expected maximum latency [37].

Price of Anarchy

For weighted congestion games on m parallel links, mixed equilibria, and social cost defined as the expected maximum latency, the price of anarchy is $\Theta(\frac{\log m}{\log \log m})$ for identical links [26, 59] and $\Theta(\frac{\log m}{\log \log \log m})$ [26] for related links. If instead social cost is defined as the sum of private costs $\frac{n}{5}$ is a lower bound on the price of anarchy [14] for pure equilibria on identical links where n is the number of weighted players.

4.3 Congestion Games with Player-Specific Latency Functions

We gave a formal definition of congestion games with player-specific latency functions in Section 3.1. Milchtaich [69] showed that there is an unweighted congestion game with player-specific latency functions on parallel links that does not have the finite best-reply property. Nevertheless, a pure Nash equilibrium for such a game is guaranteed to exist and can be computed in polynomial time [69]. The existence of pure Nash equilibria is even guaranteed for unweighted congestion games with player-specific latency functions where the strategy set of each player consists of the bases of a matroid on the set of resources [2]. Furthermore, equilibria existence and the finite improvement property are guaranteed for all unweighed congestion games with player-specific additive constants [31].

For weighted congestion games with player-specific latency functions on parallel links, there is a counterexample to the existence of a pure Nash equilibrium with only 3 players and 3 links [69]. This is a tight result since such games possess the finite best-reply property in case of 2 players and the finite improvement property in case of 2 links [69]. An explicit potential function for the latter class of games was given by Anantharam [7].

The special case of congestion games with player-specific linear latency functions on parallel links was studied by Georgiou et al. [48]. For the case of 3 weighted players, such games are guaranteed to have a pure Nash equilibrium [48]. For the case of 2 links and weighted players, there is a polynomial time algorithm to compute a pure Nash equilibrium [48]. Georgiou et al. [48] also proved upper bounds on the price of anarchy for both social cost defined as the maximum private cost of a player and social cost defined as the sum over the private costs of all players.

4.4 Wardrop Games

In recent years, Wardrop games — which were already introduced in the 1950's (see, e.g., [12, 92]) — received a lot of attention. Since a Wardrop equilibrium is a solution to a convex program, it can be computed in polynomial time using the ellipsoid method of Khachiyan [56]. This results also implies that the total latency is the same for all Wardrop equilibria. A recent paper by Fischer et al. [36] shows that a fast convergence to a Wardrop equilibrium is possible with a replication and exploration rerouting policy if a symmetric Wardrop game with polynomial latency functions is considered.

Roughgarden and Tardos analyzed the price of anarchy for Wardrop games where the total latency measures social cost. They showed that the price of anarchy is $\frac{4}{3}$ for affine latency functions [85] and $\Theta(\frac{d}{\ln d})$ for polynomials of degree at most d with non-negative coefficients [84]. If all latency functions are linear then every Wardrop equilibrium has optimum social cost [85]. Roughgarden [82] proved that the price of anarchy is independent of the network topology if a class $\mathcal{F}$ of latency functions is considered that only fulfills relatively weak assumptions. Instead, it only depends on the so called "anarchy value" $\alpha(\mathcal{F})$ of $\mathcal{F}$, and the worst-case ratio is already achieved on parallel links.

Roughgarden [84] also considered Wardrop games on networks with M/M/1 latency functions. When r is the amount of traffic and $c_{\min} > r$ is the minimum capacity among all edge capacities in the network, an upper bound on the price of anarchy is given by $\frac{1}{2} \cdot (1 + \sqrt{c_{\min}/(c_{\min} - r)})$. Observe that this expression approaches ∞ as the amount of traffic r approaches $c_{\min}$. This upper bound is asymptotically tight even for games on so-called union of paths graphs, i.e., on graphs that consist of disjoint paths from an origin node v_o to a destination node v_d where the paths only have the two nodes v_o and v_d in common.

There is also some work that focused on the price of anarchy for Wardrop games but did not use total latency to measure the social cost [24, 25, 83].

4.5 Bottleneck Games

In a recent paper, Cole et al. [22] studied Wardrop-like games where the latency of a path is defined as the p-norm, $1 < p \leq \infty$, of the vector of its edge latencies. In this context, they also looked at the case of "elastic traffic" where some share of the participants might be better off by not traveling at all.

When $p = \infty$ and in the case of inelastic traffic, their games are equal to the bottleneck games with splittable traffic that we introduced in Section 3.4. However, they looked at a subclass of Wardrop equilibria that they define as "subpath-optimal", with the reason for their restricting being that otherwise the price of anarchy is infinite even if latency functions are just linear. They showed that the anarchy value [82] is an upper bound on the price of anarchy for subpath optimal equilibria and hence also an upper bound on the price of stability.

There are also two recent papers [18, 19] investigating the price of anarchy for bottleneck congestion games where the private cost of a player is defined as the maximum latency of any edge on the path used by the player.

4.6 Finite Splittable Routing Games

There are several papers (see e.g. [6, 20, 23, 53, 58, 74]) studying finite splittable routing games where a finite number of players with non-negligible effect on each other is given. The players have to split their traffics over the available paths with the objective to minimize their private costs. In this setting, the price of anarchy for affine latency functions and social cost as total latency is at most $\frac{3}{2}$ [23].

Two papers [58, 74] studied such games with certain player-specific private cost functions that are based on M/M/1 latency functions. Korilis et al. [58] studied what happens to the private costs of the players if new capacity is added to the network or if existing capacity is reallocated. Orda et al. [74] considered the (non-)uniqueness of equilibria.

Banner and Orda [11] studied finite splittable routing games where the private cost of a player is defined as the maximum among all latencies of edges to which this player assigns a non-zero flow, whereas social cost is given by the maximum edge latency in the network. Moreover, they proved the existence and non-uniqueness of equilibria. They were also able to show that the price of anarchy is unbounded.

4.7 Harsanyi's Bayesian Games

The Nobel laureate Harsanyi developed in his pioneering work [51, 52] a framework for studying competitive situations where the players have incomplete information. For an introduction to these so-called Bayesian games, we refer to [41, 65, 71]. We will also briefly sketch the concept of Bayesian games in Section 8.1.

A class of Bayesian games related to this thesis was introduced by Facchini et al. [31]. They extended unweighted congestion games to an incomplete information setting in which each agent has, according to his type, specific rewards for the subsets of resources that he can use. Facchini et al. [31] proved that in this setting equilibria are guaranteed to exist.

Beier et al. [13] focused on a service provider game with incomplete information. In such a game, the decision of a customer to join or refuse a service does not only depend on the quality of service but also on the customer's type that describes the customer's minimum quality of service requirement.

CHAPTER 5

Congestion Games with Player-Specific Constants

5.1 Introduction

In this chapter, we study the *congestion games with player-specific constants* that we defined in Section 3.1. In a congestion games with player-specific constants each player-specific latency function $f_{ie}(x) = g_e(x) \odot c_{ie}$ is made up of a resource-specific *delay function* g_e and a player-specific *constant* c_{ie} (for the particular resource); the two are composed by means of a binary operation $\odot$. For example, *congestion games with player-specific additive constants* (resp., *multiplicative constants*) correspond to the case where the binary operation is addition (resp., multiplication).

Note that this new model of congestion games restricts Milchtaich's one [69] since player-specific latency functions are no longer completely arbitrary; simultaneously, it generalizes Rosenthal's model [80] since it allows composing player-specific constants into each latency function. We will observe that the class of congestion games with player-specific constants is contained in the more general class of *dominance congestion games* that we also defined in Section 3.1. In this more general class of congestion games with player-specific latency functions, it holds that for any pair of players, the preference of some of the two players with regard to a pair of resources necessarily induces an identical preference for the other player.

5.1.1 Contribution

We focus on pure Nash equilibria for congestion games with player-specific constants; for these equilibria, we study questions of existence, computational complexity and convergence via selfish steps (finite improvement property) and greedy selfish steps (finite best-reply property) of players. Our findings are as follows:

- Games on parallel links:
  - Every unweighted congestion game with player-specific constants on parallel links has a potential function; hence, it has the finite improvement property and a pure Nash equilibrium.
  - There is a weighted congestion game with player-specific additive constants and 3 players on 3 parallel links that does not have the finite best-reply property (and hence neither the finite improvement property).
  - There is a particular greedy selfish step cycle for weighted congestion games with general player-specific latency functions and 3 players on parallel links whose outlaw implies the existence of a pure Nash equilibrium. This cycle is indeed outlawed for weighted dominance congestion games with 3 players on parallel links – and hence for weighted congestion games with player-specific constants and 3 players on parallel links. Hence, weighted congestion games with player-specific constants and 3 players on parallel links have a pure Nash equilibrium.
- Network games:
  For unweighted symmetric network congestion games with player-specific additive constants, it is PLS-complete to find a pure Nash equilibrium.
- Arbitrary (non-network) games:
  Every weighted congestion game with linear delay functions and player-specific additive constants has a potential function; hence, it has the finite improvement property and a pure Nash equilibrium.

5.1.2 Related Work

The related work that is relevant for this chapter studies congestion games with player-specific latency functions (Section 4.3), the existence and computation of Nash equilibria in congestion games (Section 4.1) also for the special case of congestion games on parallel links (Section 4.2).

5.1.3 Road Map

We partition the results of this chapter according to the structure of the strategy sets in the congestion games with player-specific constants. We restrict to games on parallel links in Section 5.2, to games on networks in Section 5.3, whereas there is no restriction on the strategy sets in Section 5.4. In Section 5.5 we conclude and discuss directions for further research.

5.2 Congestion Games on Parallel Links

We will now show that every unweighted congestion game with player-specific constants on parallel links has the finite improvement property and a pure Nash equilibrium. For the proof we introduce a function Φ with

$$\Phi(\mathsf{s}) \;=\; \bigodot_{e\in E} \bigodot_{i=1}^{\delta_e(\mathsf{s})} g_e(i) \odot \bigodot_{i=1}^{n} c_{is_i}.$$

for any strategy profile s. We now prove that this function is a potential function:

Theorem 5.1: *Every unweighted congestion game with player-specific constants on parallel links has a potential function.*

Proof: Fix a strategy profile s. Consider a selfish step of player $k \in [n]$ to strategy t_k, which transforms s to t. Clearly, $\mathsf{PC}_k(\mathsf{s}) > \mathsf{PC}_k(\mathsf{t})$ or

$$g_{s_k}(\delta_{s_k}(\mathsf{s})) \odot c_{ks_k} \;>\; g_{t_k}(\delta_{t_k}(\mathsf{t})) \odot c_{kt_k}. \tag{5.1}$$

Note also that $\delta_{s_k}(\mathsf{t}) = \delta_{s_k}(\mathsf{s}) - 1$ and $\delta_{t_k}(\mathsf{t}) = \delta_{t_k}(\mathsf{s}) + 1$, while $\delta_e(\mathsf{t}) = \delta_e(\mathsf{s})$ for all $e \in E \setminus \{s_k, t_k\}$. Hence,

$$\begin{aligned}
\Phi(\mathsf{s}) \;&=\; \bigodot_{\substack{e\in E\setminus\\ \{s_k,t_k\}}} \bigodot_{i=1}^{\delta_e(\mathsf{s})} g_e(i) \odot \bigodot_{i\in[n]\setminus\{k\}} c_{is_i} \odot \bigodot_{i=1}^{\delta_{s_k}(\mathsf{s})} g_{s_k}(i) \odot \bigodot_{i=1}^{\delta_{t_k}(\mathsf{s})} g_{t_k}(i) \odot c_{ks_k} \\
&=\; \bigodot_{\substack{e\in E\setminus\\ \{s_k,t_k\}}} \bigodot_{i=1}^{\delta_e(\mathsf{s})} g_e(i) \odot \bigodot_{i\in[n]\setminus\{k\}} c_{is_i} \odot \bigodot_{i=1}^{\delta_{s_k}(\mathsf{s})-1} g_{s_k}(i) \\
&\quad \odot \bigodot_{i=1}^{\delta_{t_k}(\mathsf{s})} g_{t_k}(i) \odot g_{s_k}(\delta_{s_k}(\mathsf{s})) \odot c_{ks_k} \\
&\overset{(5.1)}{>} \bigodot_{\substack{e\in E\setminus\\ \{s_k,t_k\}}} \bigodot_{i=1}^{\delta_e(\mathsf{s})} g_e(i) \odot \bigodot_{i\in[n]\setminus\{k\}} c_{is_i} \odot \bigodot_{i=1}^{\delta_{s_k}(\mathsf{s})-1} g_{s_k}(i) \\
&\quad \odot \bigodot_{i=1}^{\delta_{t_k}(\mathsf{s})} g_{t_k}(i) \odot g_{t_k}(\delta_{t_k}(\mathsf{t})) \odot c_{kt_k} \\
&=\; \bigodot_{\substack{e\in E\setminus\\ \{s_k,t_k\}}} \bigodot_{i=1}^{\delta_e(\mathsf{t})} g_e(i) \odot \bigodot_{i\in[n]\setminus\{k\}} c_{is_i} \odot \bigodot_{i=1}^{\delta_{s_k}(\mathsf{t})} g_{s_k}(i) \odot \bigodot_{i=1}^{\delta_{t_k}(\mathsf{t})} g_{t_k}(i) \odot c_{kt_k} \\
&=\; \Phi(\mathsf{t}),
\end{aligned}$$

so that Φ is a potential function. ■

Theorem 5.1 immediately implies:

Corollary 5.1: *Every unweighted congestion game with player-specific constants on parallel links has the finite improvement property and a pure Nash equilibrium.*

We would like to give two remarks on the results of Theorem 5.1 and Corollary 5.1:

- Note that the theorem and the corollary even hold if the functions g_e are not non-decreasing.
- "$\odot$" is "+": Facchini et al. [31] showed a more general result for additive player-specific constants that does not restrict to parallel links. They proved the finite improvement property for all unweighted congestion games with additive player-specific constants.

We will now show that it is not possible to generalize Theorem 5.1 to the setting that allows players of different weights. To do so, we give a game with weighted players that neither possesses the finite improvement property nor the finite best-reply property.

Theorem 5.2: *There is a weighted congestion game with additive player-specific constants and 3 players on 3 parallel links that does* not *have the finite best-reply property.*

Proof: By construction. The weights of the 3 players are $w_1 = 2$, $w_2 = 1$, and $w_3 = 1$. For a load x on link e, player i's latency on this link e is given by $c_{ie} + g_e(x)$, where:

c_{ie}	Link 1	Link 2	Link 3
Player 1	0	∞	5
Player 2	0	0	∞
Player 3	∞	0	2

	Link 1	Link 2	Link 3
$g_e(1)$	1	2	1
$g_e(2)$	8	13	2
$g_e(3)$	14	∞	10

Notice that the strategy profiles $(1, 2, 3)$ and $(3, 1, 2)$ are Nash equilibria. Consider now the cycle $(1, 1, 3) \to (1, 1, 2) \to (1, 2, 2) \to (3, 2, 2) \to (3, 2, 3) \to (3, 1, 3) \to (1, 1, 3)$. The private cost of the deviating player decreases in each of these steps:

	PC_1	PC_2	PC_3
$(1, 1, 3)$	14		3
$(1, 1, 2)$		14	2
$(1, 2, 2)$	8	13	
$(3, 2, 2)$	7		13
$(3, 2, 3)$		2	12
$(3, 1, 3)$	15	1	

So, this is a cycle of selfish steps. Furthermore, each of the selfish steps from a link o to a link p of a player i is a *greedy* selfish step. If the player i would instead move to link q with $q \neq o$ and $q \neq p$ then his private cost would increase since $c_{iq} = \infty$. Therefore, this is a greedy selfish step cycle that can be used to construct an infinite sequence of greedy selfish steps. ■

Note that Theorem 5.2 does not outlaw the possibility that every weighted congestion game with player-specific constants on parallel links has a pure Nash equilibrium. Although we do not know whether equilibria are guaranteed to exist for the general case with an arbitrary number of players we will be able to prove that all such games with three players are guaranteed to possess a pure Nash equilibrium. To get this result we first establish that there is a particular greedy selfish step cycle whose outlaw implies the existence of a pure Nash equilibrium:

Theorem 5.3: *Let Γ be a weighted congestion game with player-specific latency functions and 3 players on parallel links. If Γ does* not *have a greedy selfish step cycle*

$$(l, j, j) \to (l, l, j) \to (k, l, j) \to (k, l, l) \to (k, j, l) \to (l, j, l) \to (l, j, j)$$

(where $l \neq j$, $j \neq k$, $l \neq k$ are any three links and $w_1 \geq w_2 \geq w_3$), then Γ has a pure Nash equilibrium.

Proof: Assume that Γ does not have a greedy selfish step cycle of the given form. We will construct a pure Nash equilibrium for Γ. We start by assigning player 1 to a best link: one that minimizes his private cost. Then, we assign player 2 to his best link (given the assignment of player 1). We now distinguish three different cases.

Case (A) Players 1 and 2 are assigned to the same link a and player 1 remains satisfied.

Case (B) Players 1 and 2 are assigned to the same link a and player 1 is now unsatisfied.

Case (C) Players 1 and 2 are assigned to different links a and b, respectively.

We will show how, in each of these cases, a pure Nash equilibrium can be reached by assigning player 3 and taking some greedy selfish steps.

Case (A): Assign now player 3 to his best link (given the assignments of players 1 and 2). If this link is different from a, we reached a Nash equilibrium. If all players are assigned to link a but the current strategy profile is not a Nash equilibrium, at least one of the players 1 and 2 is unsatisfied. We reach a Nash equilibrium by a greedy selfish step for one unsatisfied player.

Case (B): We now do a greedy selfish step for player 1. Let b, $b \neq a$, be the link to that player 1 is now assigned. Both players 1 and 2 are satisfied in the current strategy profile. Assign now player 3 to his best link.

- If this link is different from both a and b, we reached a Nash equilibrium.
- If player 3 is together with player 2, we also reached a Nash equilibrium since $w_1 \geq w_3$ and a was a best link for player 2 initially.
- There remains the case where player 3 is together with player 1. If player 1 is satisfied, we reached a Nash equilibrium. Else, we take a greedy selfish step for player 1. Now player 1 is either assigned to link a or to a link that is different from both a and b. In both cases, players 1 and 3 are obviously satisfied and player 2 is also satisfied since a was a best link for player 2 initially (even with player 1 on it). So, we reached a Nash equilibrium.

Case (C): Note that both players are satisfied in the current strategy profile. We now assign player 3 to his best link. If this link is different from both a and b, we reached a Nash equilibrium. We will now consider the two remaining cases (C1) where player 3 is assigned to a together with player 1 and (C2) where player 3 is assigned to link b together with player 2. Both cases (C1) and (C2) are shown in diagrammatic form in Figure 5.1.

Case (C1): If player 1 is satisfied, we reached a Nash equilibrium. Else, we take a greedy selfish step for player 1. Now player 1 is either assigned to link b or to some other link different from a. In both cases, players 1 and 3 are obviously satisfied. If player 2 is satisfied we reached a Nash equilibrium. Otherwise we do a greedy selfish step for player 2.

If player 1 is assigned to link b, the greedy selfish step takes player 2 either to link a or to some other link different from b. In both cases, players 1 and 2 are obviously satisfied. Player 3 is satisfied since a was a best link for player 3 initially (even with player 1 on it). So, we reached a Nash equilibrium.

Thus, the case remains where player 1 is assigned to a link c different from both a and b. The initial decision of player 2 for his best link implies that his current greedy selfish step will assign him to link a together with player 3. If the current strategy profile is not a Nash equilibrium, then player 1 is either satisfied of unsatisfied. We proceed by case analysis.

- If player 1 is satisfied, we take a greedy selfish step for player 3. The initial decision of player 3 for his best link implies that player 3 will go to link b, and both 2 and 3 are obviously satisfied. Player 1 is also satisfied since he neither can improve by switching to link a (due to his earlier greedy selfish step) nor to some other different

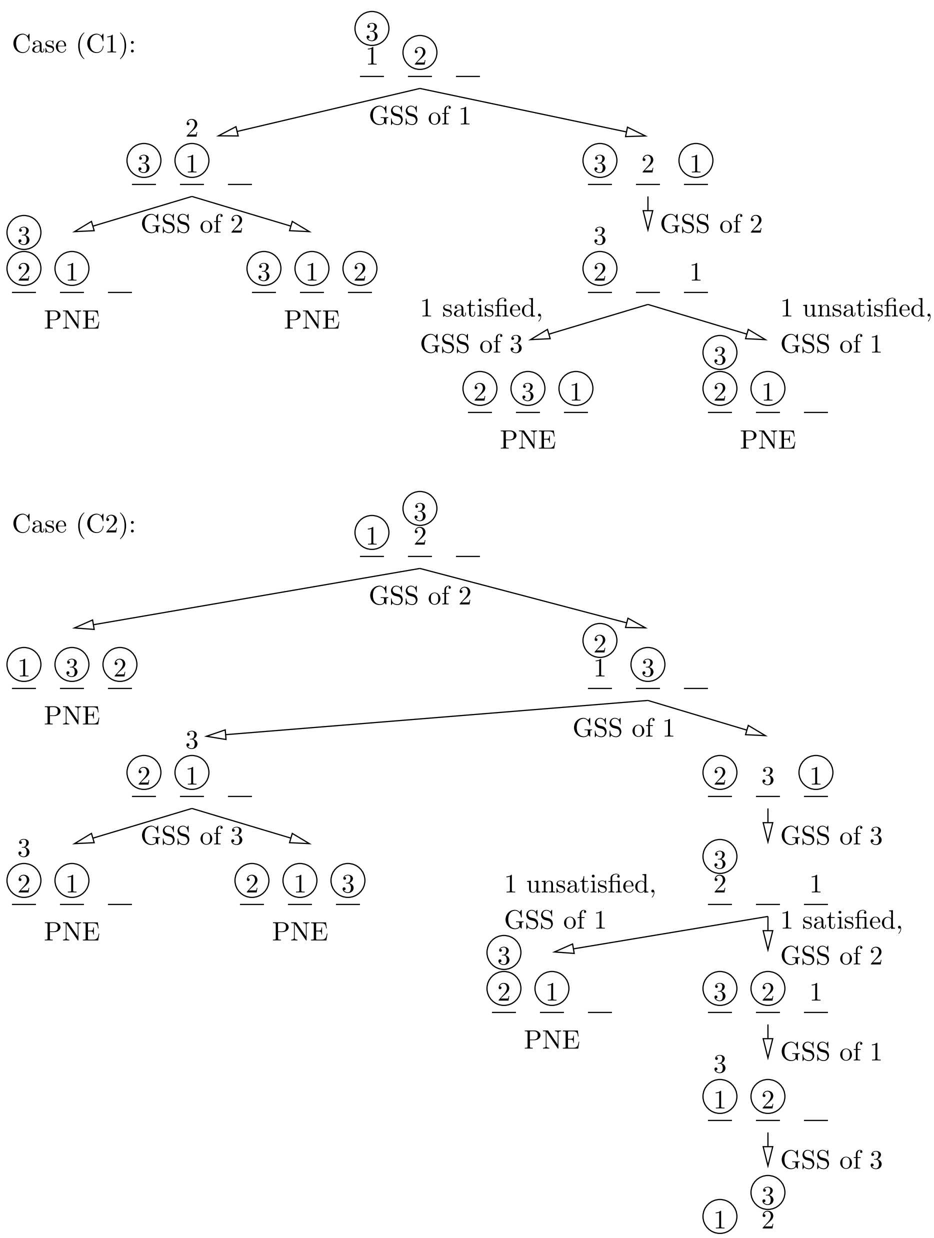

Figure 5.1: Two diagrams for the cases (C1) and (C2) that are considered in the proof of Theorem 5.3. Circles show the players that are necessarily satisfied in a given strategy profile.

link (since he was satisfied on link c before the greedy selfish step of player 3). So, we reached a Nash equilibrium.

- If player 1 is unsatisfied, we take a greedy selfish step for player 1. Now player 1 will is assigned to link b (since choosing any other link would imply a contradiction to his earlier greedy selfish step). The initial decision of player 3 for a best link implies that he is still satisfied. Player 2 is also satisfied since he neither wants to deviate to link b (due to his last greedy selfish step) nor to any other link different from a. (For the latter, observe that his initial decision for a best link implies that a deviation to a link different from a and b would induce a private cost greater than or equal to his private cost after his deviation to link b.) Since all players are satisfied, we reached a Nash equilibrium.

Case (C2): If player 2 is satisfied, we reached a Nash equilibrium. Else, we take a greedy selfish step for player 2. Now player 2 is assigned to link a or to a link different from a and b. In the latter case, we obviously reached a Nash equilibrium. If player 2 is assigned to link a, players 2 and 3 are obviously satisfied. If player 1 is also satisfied, we reached a Nash equilibrium. Else, we take a greedy selfish step for player 1 by which he is assigned either to link b or to a link c different from a and b. In both cases, both players 1 and 2 are obviously satisfied. If player 3 is also satisfied, we reached a Nash equilibrium. Else, we take a greedy selfish step for player 3.

- If player 1 is assigned to link b, then player 3 (after his greedy selfish step) either is assigned to link a or to a link different from a and b. In both cases, both players 1 and 3 are obviously satisfied; it follows from player 2's earlier greedy selfish step that he is also satisfied. So, we reached a Nash equilibrium.
- If player 1 is assigned to link c, then player 3 is (after his greedy selfish step) necessarily assigned to link a. (Player 3's initial decision for his best link implies that he cannot switch to another link.) If the current strategy profile is not a Nash equilibrium, player 1 is either unsatisfied or satisfied.

  If he is unsatisfied, we take a greedy selfish step for player 1 by which he will be necessarily assigned to link b. All other links would imply a contradiction to his earlier greedy selfish step. Player 3 is satisfied since he wants to deviate neither to link b (due to his last greedy selfish step) nor to any link other than b (due to his last greedy selfish step and his initial decision for a best link). Player 2's last greedy selfish step implies that he is also satisfied. So, we have reached a Nash equilibrium.

If player 1 is satisfied, we take a greedy selfish step for player 2 by which he is necessarily assigned to link b. All other links would imply a contradiction to his initial decision for a best link. Note that both players 2 and 3 are now satisfied. If we have not yet reached a Nash equilibrium, we take a greedy selfish step for player 1 by which he will be necessarily assigned to link a (due to his latest greedy selfish step). Player 2 is now satisfied since he neither wants to go to link a (due to his last greedy selfish step) nor to any link other than a (due to his initial decision for a best link). If we have not yet reached a Nash equilibrium, we take a greedy selfish step for player 3. Player 3 is now necessarily assigned to link b (since all other links would imply a contradiction to his initial decision for a best link). But this would complete a greedy selfish step cycle $(a, b, b) \rightarrow (a, a, b) \rightarrow (c, a, b) \rightarrow (c, a, a) \rightarrow (c, b, a) \rightarrow (a, b, a) \rightarrow (a, b, b)$. A contradiction.

It follows that Γ has a pure Nash equilibrium. ∎

We remark that Milchtaich [69, Section 8] showed that there is a weighted congestion game on parallel links with 3 players and player-specific latency functions that possesses a greedy selfish step cycle of the form described in Theorem 5.3. Nevertheless, we now show that such a cycle can not appear if we consider the more specific class of weighted dominance congestion games.

Theorem 5.4: *Every weighted dominance congestion game with 3 players on parallel links does not have a selfish step cycle of the form*

$$(l, j, j) \rightarrow (l, l, j) \rightarrow (k, l, j) \rightarrow (k, l, l) \rightarrow (k, j, l) \rightarrow (l, j, l) \rightarrow (l, j, j)$$

where $l \neq j$, $j \neq k$, $l \neq k$ *are any three links and* $w_1 \geq w_2 \geq w_3$.

Proof: Assume, by way of contradiction, that there is a dominance congestion game with such a cycle. Since all steps in the cycle are selfish steps, one gets for player 2 that

$$f_{2j}(w_2 + w_3) > f_{2l}(w_1 + w_2), \quad (5.2)$$
$$f_{2l}(w_2 + w_3) > f_{2j}(w_2). \quad (5.3)$$

In the same way, one gets for player 3 that

$$f_{3j}(w_3) > f_{3l}(w_2 + w_3), \quad (5.4)$$
$$f_{3l}(w_1 + w_3) > f_{3j}(w_2 + w_3). \quad (5.5)$$

We proceed by case analysis on whether 2 dominates 3 for (j, l) or 3 dominates 2 for (j, l).

- Assume first that 2 dominates 3 for (j, l). Then (5.2) implies that
$$f_{3j}(w_2 + w_3) \ > \ f_{3l}(w_1 + w_2) \ \geq \ f_{3l}(w_1 + w_3)$$
(since f_{3l} is non-decreasing and $w_2 \geq w_3$), a contradiction to (5.5).

- Assume now that 3 dominates 2 for (j, l). Then, (5.4) implies that
$$f_{2l}(w_2 + w_3) \ < \ f_{2j}(w_3) \ \leq \ f_{2j}(w_2)$$
(since f_{2j} is non-decreasing and $w_2 \geq w_3$), a contradiction to (5.3).

The proof is now complete. ■

Theorems 5.3 and 5.4 immediately imply:

Corollary 5.2: *Every weighted dominance congestion game with 3 players on parallel links has a pure Nash equilibrium.*

We now prove that congestion games with player-specific constants are contained within the class of dominance congestion games.

Proposition 5.1: *Each weighted congestion game with player-specific constants is a weighted dominance congestion game.*

Proof: For a given weighted congestion game with player-specific constants, fix a pair of players $i \neq j$ and a pair of resources $e \neq e'$. We proceed by case analysis.

- Assume first that $c_{ie} \odot c_{je'} \geq c_{ie'} \odot c_{je}$. We will show that j dominates i for (e, e'). Fix a pair of numbers $x, y \in \mathbb{R}^+$. Assume that $f_{je}(x) > f_{je'}(y)$ or $c_{je} \odot g_e(x) > c_{je'} \odot g_{e'}(y)$. By translation-invariance, it follows that $c_{ie} \odot c_{je} \odot g_e(x) > c_{ie} \odot c_{je'} \odot g_{e'}(y)$. The assumption that $c_{ie} \odot c_{je'} \geq c_{ie'} \odot c_{je}$ implies that $c_{ie} \odot c_{je'} \odot g_{e'}(y) \geq c_{ie'} \odot c_{je} \odot g_{e'}(y)$. It follows that $c_{ie} \odot g_e(x) > c_{ie'} \odot g_{e'}(y)$ or $f_{ie}(x) > f_{ie'}(y)$. Hence, j dominates i for (e, e').

- Assume now that $c_{ie'} \odot c_{je} > c_{ie} \odot c_{je'}$. We will show that i dominates j for (e, e'). Fix a pair of numbers $x, y \in \mathbb{R}^+$. Assume that $f_{ie}(x) > f_{ie'}(y)$ or $c_{ie} \odot g_e(x) > c_{ie'} \odot g_{e'}(y)$. By translation-invariance, it follows that $c_{je} \odot c_{ie} \odot g_e(x) > c_{je} \odot c_{ie'} \odot g_{e'}(y)$. The assumption that $c_{ie'} \odot c_{je} > c_{ie} \odot c_{je'}$ implies that $c_{je} \odot c_{ie'} \odot g_{e'}(y) > c_{je'} \odot c_{ie} \odot g_{e'}(y)$. It follows that $c_{je} \odot g_e(x) > c_{je'} \odot g_{e'}(y)$ or $f_{je}(x) > f_{je'}(y)$. Hence, i dominates j for (e, e').

The proof is now complete. ■

By Proposition 5.1, Corollary 5.2 immediately implies:

Corollary 5.3: *Every weighted congestion game with player-specific constants and 3 players on parallel links has a pure Nash equilibrium.*

We remark that this corollary broadens the earlier result by Georgiou et al. [48, Lemma B.1] for congestion games with player-specific multiplicative constants and identity delay functions where $f_{ie}(x) = c_{ie} \cdot x$.

5.3 Network Congestion Games

Recall that every unweighted congestion game with player-specific additive constants has a pure Nash equilibrium (see the earlier remark following Corollary 5.1). Nevertheless, we establish in this section that it is PLS-complete[1] to compute one even for a game on a symmetric network although a polynomial time computation is possible for such games on parallel links [69]. We prove:

Theorem 5.5: *It is PLS-complete to compute a pure Nash equilibrium in an unweighted symmetric network congestion game with player-specific additive constants.*

Proof: Clearly, the problem of computing a pure Nash equilibrium in an unweighted symmetric congestion game with player-specific additive constants is a PLS-problem. (The set of feasible solutions is the set of all strategy profiles and the neighborhood of a strategy profile is the set of strategy profiles that differ in the strategy of exactly one player; the objective function is the potential function since a local optimum of this function is a Nash equilibrium.) To prove PLS-hardness, we use a reduction from the PLS-complete problem of computing a pure Nash equilibrium for an unweighted asymmetric network congestion game [30]. For the reduction, we construct the two appropriate functions F_1 and F_2:

The function F_1: Given an unweighted asymmetric network congestion game Γ on a network G, where $(a_i, b_i)_{i\in[n]}$ are the origin and destination nodes of the n players and $(f_e)_{e\in E}$ are the latency functions, F_1 constructs a symmetric network congestion game Γ' with n players on a graph G', as follows:

- G' includes G, where for each edge e of G, $g'_e := f_e$ and $c'_{ie} := 0$ for each player $i \in [n]$.
- G' contains a new common origin node a' for all players and a new common destination node b' for all players; for each player $i \in [n]$, we add an edge (a', a_i) with $g'_{(a',a_i)}(x) := 0, c'_{i(a',a_i)} := 0$, and $c'_{k(a',a_i)} := \infty$ for all $k \neq i$; in addition, we add for each player $i \in [n]$ an edge (b_i, b') with $g'_{(b_i,b')}(x) := 0, c'_{i(b_i,b')} := 0$, and $c'_{k(b_i,b')} := \infty$ for all $k \neq i$.

The function F_2: Consider now a pure Nash equilibrium t for Γ'. The function F_2 maps t to a strategy profile s for Γ (which, we shall prove, is a Nash equilibrium for Γ) as follows:

- Note first that for each player $i \in [n]$, t_i is a path that includes both edges (a', a_i) and (b_i, b') (since otherwise $\mathsf{PC}_i(\mathsf{t}) = \infty$). Construct s_i from t_i by eliminating the edges (a', a_i) and (b_i, b').

[1] See Section 2.4 for a definition of PLS(-complete) problems.

It remains to prove that $\mathsf{s} = F_2(\mathsf{t})$ is a Nash equilibrium for Γ. By way of contradiction, assume otherwise. Then there is a player k that can decrease his private cost in Γ by changing his path s_k to s'_k. But then player k can decrease his private cost in Γ' by changing his path $t_k = (a', a_k), s_k, (b_k, b')$ to $t'_k = (a', a_k), s'_k, (b_k, b')$. So, t is not a Nash equilibrium. A contradiction. ∎

We remark that Theorem 5.5 also holds for games on graphs with *undirected* edges since the problem of computing a pure Nash equilibrium for an unweighted asymmetric network congestion game with undirected edges is also PLS-complete [1].

5.4 Arbitrary Congestion Games

Recall that Theorem 5.2 outlawed the possibility that every weighted congestion game with player-specific additive constants on parallel links has the finite best-reply property. Nevertheless, we establish in this section that every weighted congestion game with player-specific additive constants even has the finite improvement property for the special case of linear delay functions. For this case the player-specific latency functions are of the form $f_{ie}(x) = c_{ie} + a_e \cdot x$ and we introduce a function Φ where for any strategy profile s

$$\Phi(\mathsf{s}) = \sum_{i=1}^{n} \sum_{e \in s_i} w_i \cdot (2 \cdot c_{ie} + a_e \cdot (\delta_e(s) + w_i)).$$

For any pair of player $i \in [n]$ and resource $e \in E$,

$$\phi(\mathsf{s}, i, e) = w_i \cdot (2 \cdot c_{ie} + a_e \cdot (\delta_e(\mathsf{s}) + w_i)), \qquad \text{i.e.,} \qquad \Phi(\mathsf{s}) = \sum_{i=1}^{n} \sum_{e \in s_i} \phi(\mathsf{s}, i, e).$$

We now prove that Φ is a potential function:

Theorem 5.6: *Every weighted congestion game with player-specific additive constants and linear delay functions, i.e, $f_{ie}(x) = c_{ie} + a_e \cdot x$, has a potential function.*

Proof: Fix a strategy profile s. Consider a selfish step of player $k \in [n]$ to strategy t_k, which transforms s to t. Clearly, $\mathsf{PC}_k(\mathsf{s}) > \mathsf{PC}_k(\mathsf{t})$ or

$$\sum_{e \in s_k} (a_e \cdot \delta_e(\mathsf{s}) + c_{ke}) \quad > \quad \sum_{e \in t_k} (a_e \cdot \delta_e(\mathsf{t}) + c_{ke}).$$

This implies that

$$\sum_{e \in s_k \setminus t_k} (a_e \cdot \delta_e(\mathsf{s}) + c_{ke}) \quad > \quad \sum_{e \in t_k \setminus s_k} (a_e \cdot \delta_e(\mathsf{t}) + c_{ke}). \tag{5.6}$$

Clearly,

$$\begin{aligned}\Phi(\mathsf{s})-\Phi(\mathsf{t}) &= \sum_{i\in[n]}\left(\sum_{e\in s_i}\phi(\mathsf{s},i,e)-\sum_{e\in t_i}\phi(\mathsf{t},i,e)\right)\\ &= \sum_{e\in s_k}\phi(\mathsf{s},k,e)-\sum_{e\in t_k}\phi(\mathsf{t},k,e)\\ &\quad+\sum_{i\in[n]\setminus\{k\}}\left(\sum_{e\in s_i}\phi(\mathsf{s},i,e)-\sum_{e\in t_i}\phi(\mathsf{t},i,e)\right)\end{aligned}$$

We consider the first and the second part of this expression separately. On the one hand,

$$\begin{aligned}&\sum_{e\in s_k}\phi(\mathsf{s},k,e)-\sum_{e\in t_k}\phi(\mathsf{t},k,e)\\ &\quad= \sum_{e\in s_k\setminus t_k}\phi(\mathsf{s},k,e)-\sum_{e\in t_k\setminus s_k}\phi(\mathsf{t},k,e)\\ &\quad= \sum_{e\in s_k\setminus t_k}w_k\cdot(2c_{ke}+a_e(\delta_e(\mathsf{s})+w_k))-\sum_{e\in t_k\setminus s_k}w_k\cdot(2c_{ke}+a_e(\delta_e(\mathsf{t})+w_k)).\end{aligned}$$

On the other hand,

$$\begin{aligned}&\sum_{i\in[n]\setminus\{k\}}\left(\sum_{e\in s_i}\phi(\mathsf{s},i,e)-\sum_{e\in t_i=s_i}\phi(\mathsf{t},i,e)\right)\\ &\quad= \sum_{i\in[n]\setminus\{k\}}\sum_{e\in s_i}(\phi(\mathsf{s},i,e)-\phi(\mathsf{t},i,e))\\ &\quad= \sum_{i\in[n]\setminus\{k\}}\left(\sum_{\substack{e\in\\ s_i\cap(s_k\setminus t_k)}}(\phi(\mathsf{s},i,e)-\phi(\mathsf{t},i,e))+\sum_{\substack{e\in\\ s_i\cap(t_k\setminus s_k)}}(\phi(\mathsf{s},i,e)-\phi(\mathsf{t},i,e))\right)\\ &\quad= \sum_{e\in s_k\setminus t_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(\phi(\mathsf{s},i,e)-\phi(\mathsf{t},i,e))\\ &\qquad+\sum_{e\in t_k\setminus s_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(\phi(\mathsf{s},i,e)-\phi(\mathsf{t},i,e))\\ &\quad= \sum_{e\in s_k\setminus t_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(w_i\cdot a_e\cdot(\delta_e(\mathsf{s})-\delta_e(\mathsf{t})))\\ &\qquad+\sum_{e\in t_k\setminus s_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(w_i\cdot a_e\cdot(\delta_e(\mathsf{s})-\delta_e(\mathsf{t})))\\ &\quad= \sum_{e\in s_k\setminus t_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(w_i\cdot a_e\cdot w_k)+\sum_{e\in t_k\setminus s_k}\sum_{i\in[n]\setminus\{k\}\,|\,e\in s_i}(w_i\cdot a_e\cdot(-w_k))\\ &\quad= w_k\cdot\sum_{e\in s_k\setminus t_k}a_e\cdot(\delta_e(\mathsf{s})-w_k)-w_k\cdot\sum_{e\in t_k\setminus s_k}a_e\cdot(\delta_e(\mathsf{t})-w_k)\,.\end{aligned}$$

Putting these together yields that

$$\begin{aligned}\Phi(\mathsf{s}) - \Phi(\mathsf{t}) &= w_k \cdot \sum_{e \in s_k \setminus t_k} (2 \cdot c_{ke} + a_e \cdot (\delta_e(\mathsf{s}) + w_k) + a_e \cdot (\delta_e(\mathsf{s}) - w_k)) \\ &\quad - w_k \cdot \sum_{e \in t_k \setminus s_k} (2 \cdot c_{ke} + a_e \cdot (\delta_e(\mathsf{t}) + w_k) + a_e \cdot (\delta_e(\mathsf{t}) - w_k)) \\ &= 2 \cdot w_k \cdot \left(\sum_{e \in s_k \setminus t_k} (c_{ke} + a_e \cdot \delta_e(\mathsf{s})) - \sum_{e \in t_k \setminus s_k} (c_{ke} + a_e \cdot \delta_e(\mathsf{t})) \right) \\ &\overset{(5.6)}{>} 0,\end{aligned}$$

so that Φ is a potential function. ∎

Fotakis et al. [40] introduced a potential function for weighted asymmetric network congestion games with affine latency functions. If such a game with $c_{ie} = c_{je}$ for all players i, j and resources e is considered our function Φ reduces to their potential function.

Theorem 5.6 immediately implies:

Corollary 5.4: *Every weighted congestion game with player-specific additive constants and linear delay functions, i.e, $f_{ie}(x) = c_{ie} + a_e \cdot x$, has the finite improvement property and a pure Nash equilibrium.*

5.5 Conclusion and Directions for Further Research

In this chapter, we studied congestion games with player-specific constants. We showed that such games have the finite improvement property if they meet certain conditions. Moreover, we established that in the case of 3 players on parallel links a pure Nash equilibrium is guaranteed to exist although there is such a game that does not have the finite best-reply property. With respect to the computation of a pure Nash equilibrium we proved that this task is PLS-complete even for games on a symmetric network. Our work leaves open several interesting problems. On the most concrete level, we would like to ask:

- Since Corollary 5.3 only applies to the case of three players the question remains whether every weighted congestion game with player-specific constants on parallel links has a pure Nash equilibrium. This seems to be a challenging open problem even for the case of only four players.
- While Theorem 5.6 showed that all weighted congestion games with player-specific additive constants and linear delay functions have the finite improvement property it is still unknown whether this also holds for other interesting classes of delay functions.

One could also consider the price of anarchy for congestion games with player-specific constants or study a Wardrop-like splittable traffic setting with player-specific constants.

CHAPTER 6

Congestion Games with Player-Specific Affine Latency Functions

6.1 Introduction

A significant part of the work that has been done up to now on congestion games considers games where the latency functions are *affine*, i.e., $f_e(x) = a_e \cdot x + b_e$. Some important questions including the price of anarchy (see e.g. [9, 21, 26, 59, 60, 67]) and the characterization of Nash equilibria maximizing social cost (see e.g. [37, 43, 44, 63, 64]) turned out to be hard even for games with these simple latency functions. If congestion games on parallel links are used to model scheduling scenarios affine latency functions $f_e(x) = a_e \cdot x + b_e$ are of interest since an edge e can be used do model a processor with a processing speed $\frac{1}{a_e}$ and an initial load of b_e (see e.g. [60]).

In this chapter, we will study congestion games where the latency functions are at the same time *affine* and *player-specific*. Although Milchtaich [69] studied congestion games with player-specific latency functions he did not focus on affine but on general non-decreasing latency functions. Since he showed that there are games with non-decreasing latency functions that do not possess the *finite best-reply property* even in the case of unweighted players and parallel links, one goal of our work is to find out whether this negative result also holds for the restricted class of games with player-specific affine latency functions. In addition, we will, in contrast to Milchtaich [69], also study the *price of anarchy*.

6.1.1 Contribution

Our main contributions are the characterization of games that admit the finite improvement property and the extension of the techniques from [9, 21] to prove upper bounds on the price of anarchy also for games with player-specific latency functions.

- The new potential function that we introduced in Section 5.2 implies that every unweighted network congestion game on parallel links with player-specific linear latency functions possesses the finite improvement property. We give counterexamples to show that a slight deviation from this model yields a loss of the finite improvement property:
  - There is a game that does not possess the finite best-reply property if we allow affine latency functions.
  - There is a game that does not possess the finite improvement property if we allow as game graph a concatenation of two parallel link graphs.
  - There is a game that does not possess a pure Nash equilibrium if we allow a game graph where all paths are of length at most two.
- We prove that weighted congestion games on parallel links with player-specific linear latency functions do not posses the finite improvement property in general, even if there are only three players. Nevertheless, such games are guaranteed to possess the finite improvement property if there are only two players.
- For a weighted congestion game Γ with player-specific affine latency functions we show that the price of anarchy is bounded from above by the expression[1]

  $$\frac{1}{2} \cdot \left(\Delta(\Gamma) + 2 + \sqrt{\Delta(\Gamma) \cdot (\Delta(\Gamma) + 4)} \right).$$

  We also prove an asymptotically tight lower bound that even holds for the case of unweighted players, parallel links, and player-specific linear latency functions.

6.1.2 Related Work

The related work that is relevant for this chapter focuses on congestion games with player-specific latency functions (Section 4.3), the existence and computation of Nash equilibria as well as the price of anarchy for congestion games (Section 4.1), and the existence and computation of Nash equilibria in congestion games on parallel links (Section 4.2).

[1] The parameter $\Delta(\Gamma)$ was defined in Section 3.1.2. It specifies the maximum factor by which the slopes of different player-specific affine latency functions for the same edge differ.

6.1.3 Road Map

The rest of this chapter is organized as follows. We focus on the finite improvement property for congestion games on parallel links with player-specific linear latency functions in the Sections 6.2 and 6.3. In the former section we consider the case of unweighted players whereas the case of weighted players is studied in the latter section. We then give in Section 6.4 asymptotically tight bounds on the price of anarchy for congestion games with player-specific affine latency functions. We conclude in Section 6.5.

6.2 Games with Unweighted Players

Milchtaich [69] showed that unweighted congestion games on parallel links with player-specific latency functions do not possess the finite improvement property in general. Obviously, Theorem 5.1 implies that we achieve the finite improvement property if we restrict to player-specific multiplicative constants where $f_{ie}(x) = a_{ie} \cdot g_e(x)$ and a potential function is given by

$$\Phi(\mathsf{s}) \quad = \quad \prod_{e \in E} \prod_{i=1}^{\delta_e(\mathsf{s})} g_e(i) \cdot \prod_{i=1}^{n} a_{is_i}.$$

This potential function also applies for the case of player-specific linear latency functions $f_{ie}(x) = a_{ie} \cdot x$ where $g_e(x) = x$. We now give counterexamples to show that a slight deviation from this model yields a loss of the finite improvement property.

Theorem 6.1: *Unweighted symmetric network congestion games on a graph G with player-specific latency functions do (in general)* not *possess*

(a) the finite best-reply property if the game has 3 *players, affine latency functions, and G is a parallel links graph.*

(b) the finite improvement property if the game has 2 *players, linear latency functions, and G is a concatenation of* 2 *parallel link graphs connected in series.*

(c) a pure Nash equilibrium if the game has 3 *players, linear latency functions, and all paths in G are of length at most* 2.

Proof: By construction. We treat (a), (b), and (c) separately.

Part (a): We modify the instance of Figure 1 in [69] such that all latency functions are affine. The resulting 3 player game Γ on a parallel link graph with 3 edges e_1, e_2, e_3 has 3 unweighted players. The player-specific affine latency functions are given by:

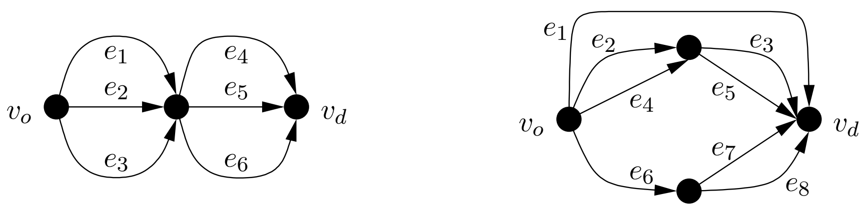

Figure 6.1: Network graphs for parts (b) (left) and (c) (right) of Theorem 6.1

$f_{i\,e}(x)$	e_1	e_2	e_3
player 1	$100 \cdot x$	$x + 10$	$10 \cdot x$
player 2	$10 \cdot x$	$100 \cdot x$	$x + 10$
player 3	$x + 10$	$10 \cdot x$	$100 \cdot x$

This game possesses a cycle of 6 greedy selfish steps $\mathsf{s}^1 \to \ldots \to \mathsf{s}^6 \to \mathsf{s}^1$:

strategy profile	s^1	s^2	s^3	s^4	s^5	s^6
strategy player 1	e_2	e_3	e_3	e_3	e_2	e_2
strategy player 2	e_1	e_1	e_3	e_3	e_3	e_1
strategy player 3	e_1	e_1	e_1	e_2	e_2	e_2
PC_1	11	10	20	20	12	12
PC_2	20	20	12	12	11	10
PC_3	12	12	11	10	20	20

Note that (e_3, e_1, e_2) is a pure Nash equilibrium for Γ.

Part (b): The network graph G with origin v_o and destination v_d is given in Figure 6.1. The player-specific linear latency functions $f_{i\,e}(x) = a_{i\,e} \cdot x$ are defined by:

$a_{i\,e}$	e_1	e_2	e_3	e_4	e_5	e_6
player 1	1	10	∞	100	100	∞
player 2	∞	100	100	∞	1	10

The game possesses a cycle of 4 selfish steps $\mathsf{s}^1 \to \ldots \to \mathsf{s}^4 \to \mathsf{s}^1$:

strategy profile	s^1	s^2	s^3	s^4
strategy player 1	$e_1\,e_5$	$e_1\,e_5$	$e_2\,e_4$	$e_2\,e_4$
strategy player 2	$e_3\,e_6$	$e_2\,e_5$	$e_2\,e_5$	$e_3\,e_6$
PC_1	$1 + 100$	$1 + 200$	$20 + 100$	$10 + 100$
PC_2	$100 + 10$	$100 + 2$	$200 + 1$	$100 + 10$

Part (c): The network graph G with origin v_o and destination v_d is given in Figure 6.1. The following table lists the values $a_{i\,e}$ for the latency functions $f_{i\,e}(x) = a_{i\,e} \cdot x$:

$a_{i\,e}$	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
player 1	150	100	∞	∞	3	∞	∞	∞
player 2	∞	∞	∞	3	100	3	∞	150
player 3	∞	3	150	∞	∞	100	3	∞

Clearly, no strategy profile that assigns a player $i \in [n]$ to a path that includes an edge e with $a_{i\,e} = \infty$ is a Nash equilibrium. Consider now the remaining strategy profiles $(\{e_1, e_2\, e_5\},\ \{e_4\, e_5, e_6\, e_8\},\ \{e_2 e_3, e_6\, e_7\})$. Six of these eight strategy profiles are no pure Nash equilibria because they belong to a cycle $\mathsf{s}^1 \to \ldots \to \mathsf{s}^6 \to \mathsf{s}^1$ of greedy selfish steps:

strategy profile	s^1	s^2	s^3	s^4	s^5	s^6
strategy pl. 1	$e_2\, e_5$	$e_2\, e_5$	$e_2\, e_5$	e_1	e_1	e_1
strategy pl. 2	$e_4\, e_5$	$e_6\, e_8$	$e_6\, e_8$	$e_6\, e_8$	$e_4\, e_5$	$e_4\, e_5$
strategy pl. 3	$e_6\, e_7$	$e_6\, e_7$	$e_2\, e_3$	$e_2\, e_3$	$e_2\, e_3$	$e_6\, e_7$
PC_1	$100+6$	$100+3$	$200+3$	150	150	150
PC_2	$3+200$	$6+150$	$3+150$	$3+150$	$3+100$	$3+100$
PC_3	$100+3$	$200+3$	$6+150$	$3+150$	$3+150$	$100+3$

We will now argue that the two remaining strategy profiles of interest are no pure Nash equilibria. Given $(e_2\, e_5, e_4\, e_5, e_2\, e_3)$ player 2 has a private cost of 3+200 and wants to deviate to $(e_2\, e_5,\ e_6\, e_8,\ e_2\, e_3)$ where his private cost is 3+150. If instead $(e_1, e_6\, e_8, e_6\, e_7)$ is given the private cost of player 3 is 200+3 and thus this player wants to deviate to $(e_1, e_6\, e_8, e_2\, e_3)$ where the private cost is $3 + 150$. ■

We would like to give two remarks on the results of Theorem 6.1.

- Part (a) of Theorem 6.1 states that there is a parallel link game with unweighted players and player-specific affine latency functions that does not possess the finite best-reply property. Note that Milchtaich [69] showed that all games of this kind possess a pure Nash equilibrium even in the more general setting allowing arbitrary non-decreasing latency functions.

- Part (b) of Theorem 6.1 states that there is a game on a concatenation of 2 parallel link graphs that does not possess the *finite improvement property*. Nevertheless, all unweighted network congestion games with player-specific linear latency functions on a graph G possess the *finite best-reply property* if G is a concatenation of parallel link graphs connected in series. To see this, observe that after a greedy selfish step of a player $i \in [n]$ this player is assigned to a best edge in each parallel link graph. The finite best-reply property follows by applying Theorem 5.1 to each parallel link graph.

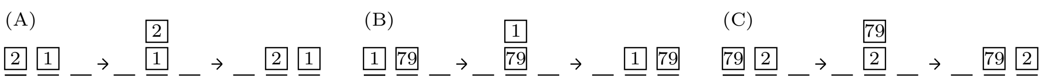

Figure 6.2: Double-steps (A), (B), (C) used in the proof of Theorem 6.2

6.3 Games with Weighted Players

6.3.1 No Finite Improvement Property for a 3 Player Game

For weighted congestion games on parallel links with player-specific linear latency functions Georgiou et al. [48] showed that a Nash equilibrium always exists in the case of 3 players. For arbitrary many players it is an open problem whether such a game still admits a pure Nash equilibrium or not. Theorem 6.2 implies that the finite improvement property can not be used to solve the open problem.

Theorem 6.2: *There is a weighted congestion game on parallel links with* 3 *players and player-specific linear latency functions that does not possess the finite improvement property.*

Proof: The 3 players of the game are of weight $w_1 = 1$, $w_2 = 2$, and $w_3 = 79$. The player-specific linear latency functions $f_{ie}(x) = a_{ie} \cdot x$ for the 11 edges are defined by:

e	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8	e_9	e_{10}	e_{11}
a_{1e}	$\frac{3^8}{80}-\epsilon_3$	∞	1	$3-\epsilon_1$	$(3-\epsilon_1)^2$	$(3-\epsilon_1)^3$	$(3-\epsilon_1)^4$	$(3-\epsilon_1)^5$	$(3-\epsilon_1)^6$	$(3-\epsilon_1)^7$	$(3-\epsilon_1)^8$
a_{2e}	∞	1	$\frac{2}{3}-\epsilon_1$	$\frac{2^2}{3^2}-\epsilon_1$	$\frac{2^3}{3^3}-\epsilon_1$	$\frac{2^4}{3^4}-\epsilon_1$	$\frac{2^5}{3^5}-\epsilon_1$	$\frac{2^6}{3^6}-\epsilon_1$	$\frac{2^7}{3^7}-\epsilon_1$	$\frac{2^8}{3^8}-\epsilon_1$	∞
a_{3e}	1	∞	$\frac{80}{79}-\epsilon_1$	∞	∞	∞	∞	∞	∞	$\frac{80^2}{79\cdot 81}-\epsilon_2$	$(\frac{80}{79}-\epsilon_1)^2$

Here the numbers ϵ_1, ϵ_2, and ϵ_3 are given by:

$$\epsilon_1 = \frac{1}{100000}, \quad \epsilon_2 = \frac{1}{10000}, \quad \epsilon_3 = 1.$$

The selfish step cycle we describe now mainly consists of double-steps (A), (B), and (C) sketched in Figure 6.2. Our cycle of selfish steps starts in the initial strategy profile (e_3, e_2, e_1). We now perform eight double-steps (A).

A double-step (A) starts in a strategy profile where player 2 is the only player who is assigned to the edge e_j. Now the first step moves player 2 to the edge e_k player 1 is assigned to. Afterwards this player 1 moves to the so far empty edge e_l. Both steps of a double-step (A) are selfish iff $a_{2e_k}/a_{2e_j} < \frac{2}{3}$ and $a_{1e_l}/a_{1e_k} < 3$. In each step of our eight double-steps (A) the deviating player moves from edge e_t to edge e_{t+1}. After the double-steps (A) the strategy profile (e_{11}, e_{10}, e_1) is reached. Observe for the first step $(e_3, e_2, e_1) \to (e_3, e_3, e_1)$ of player 2 that $a_{2e_2} \cdot \frac{2}{3} = \frac{2}{3} > \frac{2}{3} - \epsilon_1 = a_{2e_3}$.

Note for the remaining steps of player 2 in the double-steps (A) that we have for each $j \in \{3, \ldots, 9\}$ and $k = j + 1$ that

$$\begin{aligned} a_{2e_j} \cdot \frac{2}{3} &= \left(\frac{2^{j-2}}{3^{j-2}} - \epsilon_1\right) \cdot \frac{2}{3} \\ &= \frac{2^{j-1}}{3^{j-1}} - \epsilon_1 \cdot \frac{2}{3} \\ &> \frac{2^{j-1}}{3^{j-1}} - \epsilon_1 \\ &= a_{2e_{j+1}} \\ &= a_{2e_k}, \end{aligned}$$

i.e., all these steps of player 2 are selfish. For the steps that player 1 does in these double-steps (A) we get for each $k \in \{3, \ldots, 10\}$ and $l = k + 1$ that

$$\begin{aligned} a_{1e_k} \cdot 3 &= (3 - \epsilon_1)^{k-3} \cdot 3 \\ &> (3 - \epsilon_1)^{k-2} \\ &= a_{1e_{k+1}} \\ &= a_{1e_l}, \end{aligned}$$

i.e., all these steps of player 1 are selfish and hence all sixteen steps performed in the eight double-steps (A) are selfish.

The cycle continues with double-steps (B). A double-step (B) starts with a move of player 1 from edge e_j to the edge e_k used by player 3 followed by a step of player 3 to an empty edge e_l. Observe that (B) is a pair of selfish steps iff $a_{1e_k}/a_{1e_j} < \frac{1}{80}$ and $a_{3e_l}/a_{3e_k} < \frac{80}{79}$. We conduct two double-steps (B):

$$(e_{11}, e_{10}, e_1) \to (e_1, e_{10}, e_1) \to (e_1, e_{10}, e_3) \to (e_3, e_{10}, e_3) \to (e_3, e_{10}, e_{11}).$$

It is easy to see that the two steps of player 3 are selfish. For the step $(e_1, e_{10}, e_1) \to (e_1, e_{10}, e_3)$ observe that:

$$a_{3e_1} \cdot \frac{80}{79} = \frac{80}{79} > \frac{80}{79} - \epsilon_1 = a_{3e_3}.$$

Note for the step $(e_3, e_{10}, e_3) \to (e_3, e_{10}, e_{11})$ that:

$$a_{3e_3} \cdot \frac{80}{79} = \left(\frac{80}{79} - \epsilon_1\right) \cdot \frac{80}{79} > \left(\frac{80}{79} - \epsilon_1\right)^2 = a_{3e_{11}}.$$

We now consider the two steps of player 1. The first one $(e_{11}, e_{10}, e_1) \to (e_1, e_{10}, e_1)$ is selfish if $1 \cdot a_{1e_{11}} > 80 \cdot a_{1e_1}$:

$$1 \cdot (3 - \epsilon_1)^8 > 80 \cdot \left(\frac{3^8}{80} - \epsilon_3\right) \quad \text{i.e.} \quad \epsilon_3 > \frac{3^8}{80} - \frac{(3 - \epsilon_1)^8}{80}.$$

It is easy to verify that our selections of ϵ_1 and ϵ_3 fulfill these inequalities. The second step $(e_1, e_{10}, e_3) \to (e_3, e_{10}, e_3)$ of player 1 is selfish if $1 \cdot a_{1e_1} > 80 \cdot a_{1e_3}$:

$$1 \cdot \left(\frac{3^8}{80} - \epsilon_3\right) > 80 \cdot 1 \quad \text{i.e.} \quad \epsilon_3 < \frac{3^8}{80} - 80 = \frac{161}{80}.$$

These inequalities hold for our selection of ϵ_3.

Starting from the strategy profile (e_3, e_{10}, e_{11}) we proceed with a double-step (C) that moves player 3 to the edge e_{10} player 2 is assigned to and continues with a step of player 2 to the empty edge e_2. This double-step consists of selfish steps iff $a_{3e_{10}}/a_{3e_{11}} < \frac{79}{81}$ and $a_{2e_2}/a_{2e_{10}} < \frac{81}{2}$. The step of player 3 is selfish if $79 \cdot a_{3e_{11}} > 81 \cdot a_{3e_{10}}$:

$$79 \cdot \left(\frac{80}{79} - \epsilon_1\right)^2 > 81 \cdot \left(\frac{80^2}{79 \cdot 81} - \epsilon_2\right) \quad \text{i.e.} \quad \epsilon_2 > \frac{80^2}{79 \cdot 81} - \frac{79}{81} \cdot \left(\frac{80}{79} - \epsilon_1\right)^2.$$

It is again easy to check that our selections of ϵ_1 and ϵ_2 fulfill these inequalities. The step of player 2 is selfish if $81 \cdot a_{2e_{10}} > 2 \cdot a_{2e_2}$:

$$81 \cdot \left(\frac{2^8}{3^8} - \epsilon_1\right) > 2 \cdot 1 \quad \text{i.e.} \quad \epsilon_1 < \frac{2^8}{3^8} - \frac{2}{81} = \frac{94}{6561}.$$

These inequalities hold for our selection of ϵ_1.

The eleven double-steps explained up to now are followed by a final step that moves player 3 back to edge e_1: $(e_3, e_2, e_{10}) \to (e_3, e_2, e_1)$. It is selfish if $79 \cdot a_{3e_{10}} > 79 \cdot a_{3e_1}$:

$$79 \cdot \left(\frac{80^2}{79 \cdot 81} - \epsilon_2\right) > 79 \cdot 1 \quad \text{i.e.} \quad \epsilon_2 < \frac{80^2}{79 \cdot 81} - 1 = \frac{1}{6399}.$$

Our ϵ_2 fulfills these inequalities. Altogether we have that all steps in the cycle are selfish and thus the claim follows. ∎

6.3.2 Finite Improvement Property for 2 Player Games

For weighted congestion games on parallel links with 2 players and player-specific latency functions Milchtaich [69] showed that the *finite best-reply property* holds. The following theorem shows that we even get the *finite improvement property* if we restrict to player-specific linear latency functions.

Theorem 6.3: *Let Γ be a weighted congestion game on parallel links with 2 players and player-specific linear latency functions. Then, Γ possesses the finite improvement property.*

Proof: Define for a strategy profile $\mathsf{s} = (s_1, s_2)$ the function Φ_2 given by:

$$\Phi_2(\mathsf{s}) = a_{1\,s_1} \cdot a_{2\,s_2}^{\ln\left(\frac{w_1+w_2}{w_1}\right)/\ln\left(\frac{w_1+w_2}{w_2}\right)} \cdot \begin{cases} \frac{w_1+w_2}{w_1} & \text{if } s_1 = s_2, \\ 1 & \text{else.} \end{cases}$$

Let us now consider a selfish step from the strategy profile s to the strategy profile s' where a player moves from edge $i \in E$ to edge $j \in E$. It is easy to see that Φ_2 decreases if $\delta_i(\mathsf{s}) = \delta_j(\mathsf{s}') = w_k$ where $k \in \{1, 2\}$ is the deviating player. We will now show that the function Φ_2 also decreases in all other cases. To simplify notation define r by

$$r = \frac{\ln(\frac{w_1+w_2}{w_1})}{\ln(\frac{w_1+w_2}{w_2})}.$$

If player 1 uses a selfish step to switch to the edge player 2 is assigned to we have that $a_{1\,i} > a_{1\,j} \cdot \frac{w_1+w_2}{w_1}$ and therefore

$$\Phi_2(\mathsf{s}) = a_{1\,i} \cdot (a_{2\,j})^r > a_{1\,j} \cdot (a_{2\,j})^r \cdot \frac{w_1 + w_2}{w_1} = \Phi_2(\mathsf{s}').$$

If instead player 1 leaves with a selfish step the edge player 2 uses it is $a_{1\,i} \cdot \frac{w_1+w_2}{w_1} > a_{1\,j}$ and thus

$$\Phi_2(\mathsf{s}) = a_{1\,i} \cdot (a_{2\,i})^r \cdot \frac{w_1 + w_2}{w_1} > a_{1\,j} \cdot (a_{2\,i})^r = \Phi_2(\mathsf{s}').$$

We need a technical observation before we can consider the remaining cases. Observe that for any $x > 1$ we have that

$$\sqrt[\ln(x)]{x} = \sqrt[\ln(x)]{\mathsf{e}^{\ln(x)}} = \mathsf{e}.$$

Hence we get that

$$\sqrt[\ln(\frac{w_1+w_2}{w_2})]{\frac{w_1 + w_2}{w_2}} = \mathsf{e} = \sqrt[\ln(\frac{w_1+w_2}{w_1})]{\frac{w_1 + w_2}{w_1}}.$$

This implies that

$$\left(\frac{w_1 + w_2}{w_2}\right)^r = \frac{w_1 + w_2}{w_1}. \tag{6.1}$$

Consider now a selfish step where player 2 switches to the edge player 1 is assigned to. We have that $a_{2\,i} > a_{2\,j} \cdot \frac{w_1+w_2}{w_2}$ and therefore we get with (6.1) that $(a_{2\,i})^r > (a_{2\,j})^r \cdot \frac{w_1+w_2}{w_1}$ and hence

$$\Phi_2(\mathsf{s}) = a_{1\,j} \cdot (a_{2\,i})^r > a_{1\,j} \cdot (a_{2\,j})^r \cdot \frac{w_1 + w_2}{w_1} = \Phi_2(\mathsf{s}').$$

If instead player 2 leaves with a selfish step the edge player 1 uses it is $a_{2\,i} \cdot \frac{w_1+w_2}{w_2} > a_{2\,j}$ and thus $(a_{2\,i})^r \cdot \frac{w_1+w_2}{w_1} > (a_{2\,j})^r$ with (6.1). We get

$$\Phi_2(\mathsf{s}) = a_{1\,i} \cdot (a_{2\,i})^r \cdot \frac{w_1 + w_2}{w_1} > a_{1\,i} \cdot (a_{2\,j})^r = \Phi_2(\mathsf{s}').$$

This finishes the proof. ∎

6.4 Price of Anarchy

In this section we study the price of anarchy for weighted congestion games with player-specific affine latency functions.

6.4.1 Upper Bound for Games with Weighted Players

To prove our upper bound we use similar techniques as Christodoulou and Koutsoupias [21] and Awerbuch et al. [9]. The proof is also based on the following technical lemma.

Lemma 6.1: *For all $u, v \in \mathbb{R}_0^+$ and $c \in \mathbb{R}^+$ we have*

$$v \cdot (u+v) \quad \leq \quad \left(1 + \frac{1}{4c}\right) \cdot v^2 + c \cdot u^2.$$

Proof: For the proof of this lemma observe that

$$\begin{aligned}
\left(1+\frac{1}{4c}\right) \cdot v^2 + cu^2 - v \cdot (u+v) \quad &= \quad \frac{1}{4c} \cdot v^2 - uv + cu^2 \\
&= \quad \frac{1}{c} \cdot \left(\frac{v^2}{4} - cuv + c^2u^2\right) \\
&= \quad \frac{1}{c} \cdot \left(\frac{v}{2} - cu\right)^2 \\
&\geq \quad 0.
\end{aligned}$$

The claim follows. ■

Theorem 6.4: *Consider a weighted network congestion game Γ with player-specific affine latency functions and an associated mixed Nash equilibrium $\mathbf{Q}$. Then,*

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q}, \Gamma)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)} \quad \leq \quad \frac{1}{2} \cdot \left(\Delta(\Gamma) + 2 + \sqrt{\Delta(\Gamma) \cdot (\Delta(\Gamma)+4)}\right).$$

Proof: Let $\mathsf{o} = (o_1, \ldots, o_n)$ be a pure strategy profile with optimum social cost, i.e., $\mathsf{SC}_{\mathsf{TL}}(\mathsf{o}, \Gamma) = \mathsf{OPT}_{\mathsf{TL}}(\Gamma)$. Since $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is a Nash equilibrium, player i cannot improve by switching from strategy Q_i to o_i. Thus for any player $i \in [n]$,

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{Q}) \quad &\leq \quad \mathsf{PC}_i(\mathbf{Q}_{-i}, o_i) \\
&= \quad \sum_{\mathsf{s} \in S} q(\mathsf{s}) \cdot \left(\sum_{e \in o_i \cap s_i} f_{ie}(\delta_e(\mathsf{s})) + \sum_{e \in o_i \setminus s_i} f_{ie}(\delta_e(\mathsf{s}) + w_i) \right) \\
&\leq \quad \sum_{\mathsf{s} \in S} q(\mathsf{s}) \cdot \sum_{e \in o_i} f_{ie}(\delta_e(\mathsf{s}) + \delta_e(\mathsf{o})).
\end{aligned}$$

It follows that

$$\begin{aligned}
&\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma)\\
&= \sum_{i\in[n]} w_i \cdot \mathsf{PC}_i(\mathbf{Q})\\
&\leq \sum_{i\in[n]} w_i \cdot \mathsf{PC}_i(\mathbf{Q}_{-i}, o_i)\\
&= \sum_{i\in[n]} \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in o_i} w_i \cdot f_{ie}(\delta_e(\mathsf{s}) + \delta_e(\mathsf{o}))\\
&= \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \sum_{i, o_i \ni e} w_i \cdot [a_{ie} \cdot (\delta_e(\mathsf{s}) + \delta_e(\mathsf{o})) + b_{ie}]\\
&= \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \frac{\sum_{i,o_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{o})} \cdot \delta_e(\mathsf{o}) \cdot (\delta_e(\mathsf{s}) + \delta_e(\mathsf{o}))\\
&\quad + \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})=0}} \delta_e(\mathsf{o}) \cdot \sum_{i,o_i\ni e} a_{ie} \cdot w_i + \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \sum_{i,o_i\ni e} w_i \cdot b_{ie}.
\end{aligned}$$

By Lemma 6.1 we get for any $c \in \mathbb{R}^+$ that $\delta_e(\mathsf{o}) \cdot (\delta_e(\mathsf{s}) + \delta_e(\mathsf{o})) \leq \left(1 + \frac{1}{4c}\right) \cdot \delta_e(\mathsf{o})^2 + c \cdot \delta_e(\mathsf{s})^2$ and therfore

$$\begin{aligned}
&\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma)\\
&\leq \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \frac{\sum_{i,o_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{o})} \cdot \left[\left(1+\frac{1}{4c}\right) \cdot \delta_e(\mathsf{o})^2 + c\cdot \delta_e(\mathsf{s})^2\right]\\
&\quad + \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})=0}} \delta_e(\mathsf{o}) \cdot \sum_{i,o_i\ni e} a_{ie} \cdot w_i + \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \sum_{i,o_i\ni e} w_i \cdot b_{ie}\\
&\leq \left(1+\frac{1}{4c}\right) \cdot \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \left(\sum_{i,o_i\ni e} a_{ie}\cdot w_i\right) \cdot \delta_e(\mathsf{o})\\
&\quad + \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \sum_{i,o_i\ni e} w_i \cdot b_{ie} + c \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \frac{\sum_{i,o_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{o})} \delta_e(\mathsf{s})^2\\
&\leq \left(1+\frac{1}{4c}\right) \cdot \sum_{i\in[n]} w_i \cdot \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in o_i} (a_{ie}\cdot \delta_e(\mathsf{o}) + b_{ie})\\
&\quad + c \cdot \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \frac{\sum_{i,o_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{o})} \delta_e(\mathsf{s})^2\\
&= \left(1+\frac{1}{4c}\right) \cdot \mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma) + c \cdot \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E, \delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \frac{\sum_{i,o_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{o})} \cdot \delta_e(\mathsf{s})^2.
\end{aligned}$$

Observe that $\frac{1}{\delta_e(\mathsf{o})}\cdot\sum_{i,o_i\ni e} a_{ie}\cdot w_i$ is a weighted average slope of latency functions for edge $e \in E$. With $\frac{a_{ie}}{a_{ke}} \leq \Delta(\Gamma)$ for all $i,k \in [n]$ with $a_{ie}, a_{ke} < \infty$ it follows that $\frac{1}{\delta_e(\mathsf{o})} \cdot \sum_{i,o_i\ni e} a_{ie} \cdot w_i \leq \Delta(\Gamma) \cdot \frac{1}{\delta_e(\mathsf{s})} \cdot \sum_{i,s_i\ni e} a_{ie} \cdot w_i$. We get,

$$
\begin{aligned}
&\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma)\\
&\leq \left(1+\frac{1}{4c}\right)\mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma) + c\sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{\substack{e\in E,\delta_e(\mathsf{o})>0,\\ \delta_e(\mathsf{s})>0}} \Delta(\Gamma)\cdot \frac{\sum_{i,s_i\ni e} a_{ie}\cdot w_i}{\delta_e(\mathsf{s})}\cdot \delta_e(\mathsf{s})^2\\
&\leq \left(1+\frac{1}{4c}\right)\cdot\mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma) + c\cdot\sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in E} \Delta(\Gamma)\cdot \delta_e(\mathsf{s}) \cdot \sum_{i,s_i\ni e} a_{ie}\cdot w_i\\
&= \left(1+\frac{1}{4c}\right)\cdot\mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma) + c\cdot\Delta(\Gamma)\cdot\sum_{i\in[n]} w_i \cdot \sum_{\mathsf{s}\in S} q(\mathsf{s}) \cdot \sum_{e\in s_i} a_{ie}\cdot \delta_e(\mathsf{s})\\
&\leq \left(1+\frac{1}{4c}\right)\cdot\mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma) + c\cdot\Delta(\Gamma)\cdot\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma).
\end{aligned}
$$

Thus choosing $c = \frac{-\Delta(\Gamma)+\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}{4\cdot\Delta(\Gamma)}$ yields

$$
\begin{aligned}
&\frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{Q},\Gamma)}{\mathsf{SC}_{\mathsf{TL}}(\mathsf{o},\Gamma)}\\
&\leq \left(1+\frac{1}{4c}\right)\cdot\frac{1}{1-c\cdot\Delta(\Gamma)}\\
&= \frac{\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}{-\Delta(\Gamma)+\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}} + \frac{4}{4+\Delta(\Gamma)-\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}\\
&= \frac{4\cdot\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}{(2\Delta(\Gamma)+4)\cdot\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}-2\cdot[\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)]}\\
&= \frac{2}{(\Delta(\Gamma)+2)-\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}\\
&= \frac{2\cdot[(\Delta(\Gamma)+2)+\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}]}{[(\Delta(\Gamma)+2)-\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}]\cdot[(\Delta(\Gamma)+2)+\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}]}\\
&= \frac{\Delta(\Gamma)+2+\sqrt{\Delta(\Gamma)\cdot(\Delta(\Gamma)+4)}}{2}.
\end{aligned}
$$

Since $\mathbf{Q}$ is an arbitrary mixed Nash equilibrium the claim follows. ■

Interestingly, we get with Theorem 6.4 an upper bound of $\frac{1}{2}\cdot(3+\sqrt{5})$ in the case of $\Delta(\Gamma)=1$ which matches the exact price of anarchy for weighted congestion games [9] even though our model still allows for player-specific constants $b_{ie} \neq b_{ke}$.

6.4.2 Lower Bound for Games with Unweighted Players

The proof of the lower bound on the price of anarchy is based on the following technical lemma.

Lemma 6.2: *Let $t \in \mathbb{N}$, $\mu_1 \geq \mu_2 \geq \ldots \geq \mu_t > 0$, $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_t > 0$ where $\mu_i \in \mathbb{R}$ and $\lambda_i \in \mathbb{R}$ for all $i \in [t]$. Then,*

$$\frac{\sum_{i=1}^{t} \lambda_i \cdot \mu_i}{\sum_{i=1}^{t} \mu_i} \geq \frac{1}{t} \cdot \sum_{i=1}^{t} \lambda_i.$$

Proof: Observe that

$$\begin{aligned}
t \cdot \sum_{i=1}^{t} \lambda_i \mu_i - \sum_{i=1}^{t} \lambda_i \cdot \sum_{j=1}^{t} \mu_j &= \sum_{i=1}^{t} \lambda_i \cdot \sum_{j=1}^{t} (\mu_i - \mu_j) \\
&= \sum_{i=1}^{t-1} \sum_{j=i+1}^{t} \lambda_i (\mu_i - \mu_j) + \sum_{j=2}^{t} \sum_{i=1}^{j-1} \lambda_j (\mu_j - \mu_i) \\
&= \sum_{i=1}^{t-1} \sum_{j=i+1}^{t} \lambda_i (\mu_i - \mu_j) - \sum_{j=2}^{t} \sum_{i=1}^{j-1} \lambda_j (\mu_i - \mu_j) \\
&= \sum_{i=1}^{t-1} \sum_{j=i+1}^{t} (\lambda_i - \lambda_j) \cdot (\mu_i - \mu_j) \\
&\geq 0,
\end{aligned}$$

since each summand is non-negative. The claim follows immediately. ■

We now proceed with the asymptotically tight lower bound on the price of anarchy. Variations of the games used in the proof of the lower bound were also used in some recent papers to show lower bounds on the price of anarchy in different settings (see e.g. [10, 26, 42]).

Theorem 6.5: *For each $l \in \mathbb{N}$ and for each $\epsilon > 0$ there is an unweighted congestion game Γ on parallel links with player-specific linear latency functions and $\Delta(\Gamma) \geq l$ that possesses a pure Nash equilibrium s such that*

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\mathsf{s}, \Gamma)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)} \geq (1 - \epsilon) \cdot \Delta(\Gamma).$$

Proof: We will now give the construction of Γ that uses a variable $\Delta \in \mathbb{N}$. The game has $\Delta + 1$ classes of edges $M_1, \ldots, M_{\Delta+1}$. It is

$$|M_1| = 1 \quad \text{and} \quad |M_j| = (\Delta - 1) \cdot \prod_{i=3}^{j} (\Delta - i + 2) \ \text{ for all } j, 2 \leq j \leq \Delta + 1.$$

Furthermore, we define Δ classes of players $U_1, \dots, U_\Delta$ where the set U_i, $1 \leq i \leq \Delta$, consists of $(\Delta+1-i)\cdot|M_i|$ players. A player j in the set U_i, $1 \leq i \leq \Delta$, assigns to all edges e in M_i and M_{i+1} the slope $a_{je} = \Delta^{2-i}$ whereas he assigns slope $a_{j\hat{e}} = \infty$ to all other edges $\hat{e} \notin M_i \cup M_{i+1}$. This finishes the construction of Γ. Clearly, $\Delta = \Delta(\Gamma)$.

It is easy to see that a strategy profile with optimum social cost assigns exactly one player of U_1 to each edge in $M_1 \cup M_2$ and exactly one player of U_i, $2 \leq i \leq \Delta$, to each edge in M_{i+1}. Define the pure strategy profile s that assigns for each i, $1 \leq i \leq \Delta$, exactly $\Delta + 1 - i$ players from U_i to each edge in M_i.

We will now show that s is a Nash equilibrium. Consider an arbitrary player j from a set U_i, $1 \leq i \leq \Delta$. Since $a_{j\hat{e}} = \infty$ for each edge $\hat{e} \notin M_i \cup M_{i+1}$ player j does not want to deviate to an edge that is neither in M_i nor in M_{i+1}. Recall that $a_{je} = \Delta^{2-i}$ for each edge $e \in M_i \cup M_{i+1}$. Player j does not want to deviate to an edge in $M_i \cup M_{i+1}$ since there is a load of $\Delta - i$ on each edge in M_{i+1} and a load of $\Delta + 1 - i$ on each edge in M_i including the each to that player j is assigned. Hence, s is a Nash equilibrium.

We now define some values $A_1, A_2, \dots$ that will help us to bound the ratio $\mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma)/\mathsf{OPT}_{\mathsf{TL}}(\Gamma)$. Define

$$A_1 \;=\; \Delta^2 \quad \text{and} \quad A_j \;=\; \frac{\Delta-1}{\Delta^{j-2}} \cdot \prod_{i=1}^{j-1}(\Delta - i) \;\text{ for all } j,\, 2 \leq j \leq \Delta.$$

Observe that on the one hand we get for the social cost of the equilibrium s

$$\begin{aligned}
\mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma) &= \sum_{j=1}^{\Delta} |U_j| \cdot \Delta^{2-j} \cdot (\Delta + 1 - j) \\
&= \sum_{j=1}^{\Delta} (\Delta + 1 - j) \cdot |M_j| \cdot \Delta^{2-j} \cdot (\Delta + 1 - j) \\
&= \Delta^3 + \sum_{j=2}^{\Delta} (\Delta + 1 - j)^2 \cdot \Delta^{2-j} \cdot (\Delta - 1) \cdot \prod_{i=3}^{j} (\Delta - i + 2) \\
&= \sum_{j=1}^{\Delta} (\Delta + 1 - j) \cdot A_j.
\end{aligned}$$

On the other hand we have for the social cost of the optimum solution that

$$\begin{aligned}
\mathsf{OPT}_{\mathsf{TL}}(\Gamma) &= \sum_{j=1}^{\Delta} |U_j| \cdot \Delta^{2-j} \cdot 1 \\
&= \sum_{j=1}^{\Delta} (\Delta + 1 - j) \cdot |M_j| \cdot \Delta^{2-j}
\end{aligned}$$

$$= \Delta^2 + \sum_{j=2}^{\Delta}(\Delta+1-j)\cdot\Delta^{2-j}\cdot(\Delta-1)\cdot\prod_{i=3}^{j}(\Delta-i+2)$$
$$= \sum_{j=1}^{\Delta} A_j.$$

For our upcoming analysis of the values A_i let $k \in \mathbb{N}$ be an arbitrary number such that $k \leq \sqrt[3]{\Delta}$. We can assume that Δ is a multiple of k. Furthermore, let $i \in \mathbb{N}$ be a number such that $\frac{\Delta}{k}+1 \leq i \leq \Delta$. Then,

$$A_i = \frac{\Delta-i+1}{\Delta}\cdot A_{i-1} \leq \frac{\Delta-\frac{\Delta}{k}}{\Delta}\cdot A_{i-1} = \frac{k-1}{k}\cdot A_{i-1}. \tag{6.2}$$

We define $r = \frac{k-1}{k}$ and will now give an upper bound on $\sum_{j=1}^{\Delta} r^j$. Since $r < 1$ we have that $\sum_{j=0}^{\infty} r^j = \frac{1}{1-r}$ and thus:

$$\sum_{j=1}^{\Delta}\left(\frac{k-1}{k}\right)^j = \sum_{j=0}^{\Delta} r^j - 1 \leq \frac{1}{1-r} - 1 = k-1. \tag{6.3}$$

Combining (6.2) and (6.3) yields:

$$\sum_{i=\frac{\Delta}{k}+1}^{\Delta} A_i \overset{(6.2)}{\leq} \sum_{i=1}^{\Delta-\frac{\Delta}{k}}\left(\frac{k-1}{k}\right)^i\cdot A_{\frac{\Delta}{k}} \overset{(6.3)}{\leq} (k-1)\cdot A_{\frac{\Delta}{k}}. \tag{6.4}$$

For the lower bound on the ratio of the social cost values observe:

$$\begin{aligned}
\frac{\mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)} &= \frac{\sum_{j=1}^{\Delta}(\Delta+1-j)\cdot A_j}{\sum_{j=1}^{\Delta} A_j}\\
&\overset{(6.4)}{\geq} \frac{\sum_{j=1}^{\frac{\Delta}{k}}(\Delta+1-j)\cdot A_j}{(k-1)\cdot A_{\frac{\Delta}{k}} + \sum_{j=1}^{\frac{\Delta}{k}} A_j}\\
&= \frac{k\cdot(\Delta+1-\frac{\Delta}{k})\cdot\frac{1}{k}\cdot A_{\frac{\Delta}{k}} + \sum_{j=1}^{\frac{\Delta}{k}-1}(\Delta+1-j)\cdot A_j}{(k-1)\cdot A_{\frac{\Delta}{k}} + \sum_{j=1}^{\frac{\Delta}{k}} A_j}
\end{aligned}$$

We will now apply Lemma 6.2 with $t = \frac{\Delta}{k}+k-1$, $\mu_j = A_j$ and $\lambda_j = \Delta+1-j$ for all $j \in [\frac{\Delta}{k}-1]$, $\mu_j = A_{\frac{\Delta}{k}}$ and $\lambda_j = (\Delta+1-\frac{\Delta}{k})\cdot\frac{1}{k}$ for all j, $\frac{\Delta}{k} \leq j \leq \frac{\Delta}{k}+k-1$.

We get with this lemma:

$$\begin{aligned}
\frac{\mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)} &\geq \frac{1}{\frac{\Delta}{k}+k-1}\cdot\sum_{j=1}^{\frac{\Delta}{k}}(\Delta+1-j)\\
&= \frac{2k^2}{2k\Delta+2k^3-2k^2}\cdot\left(\sum_{i=1}^{\Delta} i-\sum_{i=1}^{\Delta-\frac{\Delta}{k}} i\right)\\
&= \frac{2k^2}{2k\Delta+2k^3-2k^2}\cdot\frac{1}{2}\cdot\left(\frac{2\Delta^2}{k}-\frac{\Delta^2}{k^2}+\frac{\Delta}{k}\right)\\
&= \Delta\cdot\frac{(2k-1)\cdot\Delta+k}{2k\Delta+2k^3-2k^2}\\
&\geq \Delta\cdot\frac{(2k-1)\cdot\Delta}{2k\Delta+2k^3}\\
&\overset{(\Delta\geq k^3)}{\geq} \Delta\cdot\frac{(2k-1)\cdot\Delta}{2k\Delta+2\Delta}\\
&= \Delta\cdot\frac{2k-1}{2k+2}.
\end{aligned}$$

The result follows by selecting k and Δ large enough. ∎

The construction in the just given proof uses a large number of players. However, the price of anarchy is unbounded even for 2 player games.

Theorem 6.6: *For every $k\in\mathbb{N}$ there is a weighted congestion game Γ^k on parallel links with 2 players and player-specific linear latency functions that possesses a pure Nash equilibrium s such that*

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\mathsf{s},\Gamma^k)}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma^k)} > k.$$

Proof: The game graph for Γ^k has 3 edges e_1, e_2, e_3 and 2 players of weight $w_1 = 2k$ and $w_2 = (2k)^2$. The player-specific linear latency functions $f_{i\,e}(x) = a_{i\,e}\cdot x$ are defined by:

$a_{i\,e}$	e_1	e_2	e_3
player 1	$(2k)^4$	$(2k)^3$	∞
player 2	∞	1	1

Consider the strategy profiles $\mathsf{s} = (e_1, e_2)$ and $\mathsf{s}' = (e_2, e_3)$. We will now argue that s is a Nash equilibrium. If player 1 moves to edge e_2 its private cost increases from $32k^5$ to $32k^5 + 16k^4$. The private cost of player 2 does not change if the player deviates to e_3. Thus s is a Nash equilibrium and we get with respect to social cost of the two strategy profiles that

$$\begin{aligned} \mathsf{SC}_{\mathsf{TL}}(\mathsf{s}, \Gamma^k) &= 64k^6 + 16k^4 \\ &> k \cdot (32k^5 + 16k^4) \\ &= k \cdot \mathsf{SC}_{\mathsf{TL}}(\mathsf{s}', \Gamma^k). \end{aligned}$$

The claim follows. ∎

6.5 Conclusion and Directions for Further Research

In this chapter, we focused on congestion games with player-specific affine latency functions. For this setting, we gave some insight into the conditions that such games have to meet in order to have the finite improvement property, the finite best-reply property, or a pure Nash equilibrium. Furthermore, we proved an upper and an asymptotically tight lower bound on the price of anarchy. There are still some open problems in this setting:

- While Theorem 6.2 showed that there is a weighted congestion game with player-specific linear latency functions on parallel links that does not possess the finite improvement property it is still unknown whether such a game is guaranteed to possess (1) a pure Nash equilibrium and (2) the finite best-reply property.
- Since we only focused on the price of anarchy for social cost as total latency the price of anarchy with respect to other social cost measures is still unknown.

CHAPTER 7

Wardrop Games with Player-Specific Affine Latency Functions

7.1 Introduction

Player-specific latency functions are reasonable for network routing games whenever the players have different beliefs, estimates, or preferences. Up to now player-specific latency functions were on the one hand considered for congestion games [31, 47, 48, 69] where the traffics of the players are *unsplittable* and in an equilibrium each player minimizes its private cost. On the other hand, such latency functions were studied for finite splittable routing games [74] where the traffics of the players are *splittable* and again in an equilibrium each player minimizes its private cost.

We focus in this chapter on the *Wardrop games with player-specific latency functions* that we defined in Section 3.2. This generalization of the original Wardrop games allows for player-specific latency functions. The traffics of the players are *splittable* and in an equilibrium each player assigns its traffic in such a way that no fraction of the traffic assigned to some path, however small, can decrease its experienced latency by unilaterally switching to another path.

7.1.1 Contribution

In this chapter, we consider the existence of equilibria, the computation of equilibria, and the price of anarchy.

- For Wardrop games on parallel links with player-specific linear latency functions, we introduce a new convex potential function and show that this function is minimized if and only if the corresponding strategy profile is a Wardrop equilibrium. This result implies that for this setting a Wardrop equilibrium can be computed in polynomial time.

- We prove that the set of equilibria of each Wardrop game on parallel links with strictly increasing player-specific latency functions is convex. But we also show that this does not hold for all games on general networks with player-specific linear latency functions. Therefore, a convex potential function does not exist for the latter setting.

- For Wardrop games Υ on arbitrary networks with player-specific affine latency functions, we show that the price of anarchy is bounded from above by the parameter[1] $\Delta(\Upsilon)$ if $\Delta(\Upsilon) > 2$ and by $\frac{4}{4-\Delta(\Upsilon)}$ otherwise. We also show a lower bound of $\frac{1}{4}\sqrt{\Delta(\Upsilon)}$ which already holds for the case of player-specific linear latency functions and parallel links.

- For Wardrop games with strictly increasing player-specific latency functions, we show that a Wardrop equilibrium always exists.

7.1.2 Related Work

The related work that is relevant for this chapter studies Wardrop games (Section 4.4) and finite splittable routing games (Section 4.6).

7.1.3 Road Map

The rest of this chapter is organized as follows. We present our results on the convergence to a Wardrop equilibrium in Section 7.2. In Section 7.3 we consider whether it is possible to extend these results to a more general setting. We also study the price of anarchy (Section 7.4) and the existence of Wardrop equilibria (Section 7.5). Finally, we conclude (Section 7.6).

7.2 Existence of and Convergence to a Wardrop Equilibrium

In this section, we consider Wardrop games with player-specific linear latency functions on parallel links. For such a game and a strategy profile $\mathbf{x}$ define the following function:

$$\Psi(\mathbf{x}) \;=\; \sum_{i\in[n]}\sum_{e\in E} x_{i\,e}\cdot\ln(a_{i\,e}) + \sum_{\substack{e\in E,\\ \delta_e(\mathbf{x})>0}} \delta_e(\mathbf{x})\cdot\ln(\delta_e(\mathbf{x})).$$

[1] The parameter $\Delta(\Upsilon)$ was defined in Section 3.2.2. It specifies the maximum factor by which the slopes of different player-specific affine latency functions for the same edge differ.

Note, that

$$\mathsf{e}^{\Psi(\mathbf{x})} = \prod_{i\in[n]}\prod_{e\in E} a_{i\,e}^{x_{i\,e}} \cdot \prod_{\substack{e\in E,\\ \delta_e(\mathbf{x})>0}} \delta_e(\mathbf{x})^{\delta_e(\mathbf{x})}$$

has a similar form as the potential function Φ used in Section 6.2. Furthermore, Ψ plays a similar role as the potential function Φ:

Theorem 7.1 (Gairing, Monien, Tiemann [47]): *Let Υ be a Wardrop game with player-specific linear latency functions on parallel links. Moreover, let $\mathbf{x}$ be a strategy profile for Υ so that there exists a player $k \in [n]$, two edges $p, q \in E$, and some Λ, $0 < \Lambda \leq x_{k\,p}$ such that:*

$$a_{k\,p} \cdot (\delta_p(\mathbf{x}) - \Lambda) \geq a_{k\,q} \cdot (\delta_q(\mathbf{x}) + \Lambda)\,.$$

Define a new strategy profile $\mathbf{y}$ by:

$$y_{i\,j} = \begin{cases} x_{k\,p} - \Lambda & \text{if } i = k, j = p,\\ x_{k\,q} + \Lambda & \text{if } i = k, j = q,\\ x_{i\,j} & \text{otherwise.}\end{cases}$$

Then $\Psi(\mathbf{y}) < \Psi(\mathbf{x})$.

This theorem can be used to show the direction (a) $\Rightarrow$ (b) for the next theorem stating that $\Psi(\mathbf{x})$ is minimized iff $\mathbf{x}$ is a Wardrop equilibrium:

Theorem 7.2 (Gairing, Monien, Tiemann [47]): *Let Υ be a Wardrop game with player-specific linear latency functions on parallel links. Moreover, let $\mathbf{y}$ be a strategy profile for Υ. Then the following two conditions are equivalent:*

(a) $\Psi(\mathbf{y}) = \min_{\mathbf{x}\in\mathcal{X}} \Psi(\mathbf{x})$,

(b) $\mathbf{y}$ *is a Wardrop equilibrium,*

where $\mathcal{X}$ is the set of all strategy profiles for the game Υ.

Since Ψ is a convex function it follows with Theorem 7.2 that the ellipsoid method of Khachiyan [56] can be used to compute a Wardrop equilibrium in polynomial time:

Theorem 7.3: *For every Wardrop game with player-specific linear latency functions on parallel links a Wardrop equilibrium can be computed in time polynomial in the size of the instance and the number of bits of precision required.*

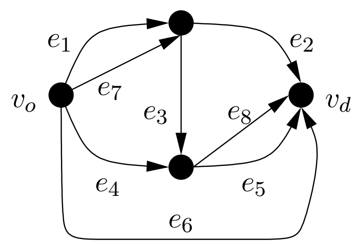

Figure 7.1: Graph used in the proof of Theorem 7.4

7.3 Is there a Convex Potential Function for a more General Setting?

If a game can be described by a convex potential function then the set of Nash equilibria forms a convex set. In this section, we show that no such convex function exists for general graphs with player-specific linear latency functions (Theorem 7.4) whereas the existence remains an open problem for parallel links with strictly increasing player-specific latency functions (Theorem 7.5).

Theorem 7.4: *There is a Wardrop game Υ with player-specific linear latency functions that possesses two Wardrop equilibria $\mathbf{x}$ and $\mathbf{y}$ where*

(a) $\delta_e(\mathbf{x}) \neq \delta_e(\mathbf{y})$ *for an edge* $e \in E$ *and* $\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}) \neq \mathsf{SC}_{\mathsf{TL}}(\mathbf{y})$,

(b) the set of Wardrop equilibria for Υ *does not form a convex set.*

Proof: The game Υ has 4 players of traffic $w_1 = w_2 = w_3 = 1$ and $w_4 = 14$. The game graph G is sketched in Figure 7.1. The player-specific linear latency functions $f_{ie}(u) = a_{ie} \cdot u$ are defined by:

a_{ie}	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
player 1	1	∞	1	∞	1	9	∞	∞
player 2	4	1	∞	1	4	2	∞	∞
player 3	∞	∞	∞	∞	∞	1	∞	∞
player 4	∞	∞	1	∞	∞	∞	1	1

The paths $\{P_2, P_3, P_4, P_6, P_7\}$ from the origin vertex v_o to the destination vertex v_d will be important in the rest of this proof. It is

$$P_2 = (e_1, e_2),\ P_3 = (e_1, e_3, e_5),\ P_4 = (e_4, e_5),\ P_6 = (e_6),\ P_7 = (e_7, e_3, e_8).$$

We define now two strategy profiles $\mathbf{x}$ and $\mathbf{y}$. $\mathbf{x}$ is given by $x_{1P_3} = x_{2P_6} = x_{3P_6} = 1$, $x_{4P_7} = 14$ whereas $\mathbf{y}$ is defined by $y_{1P_3} = \frac{1}{4}$, $y_{1P_6} = \frac{3}{4}$, $y_{2P_2} = y_{2P_4} = \frac{1}{2}$, $y_{3P_6} = 1$, and $y_{4P_7} = 14$. We will now show that both $\mathbf{x}$ and $\mathbf{y}$ are Wardrop equilibria.

- Consider the strategy profile $\mathbf{x}$. Player 1 does not want to deviate since

$$\sum_{e \in P_3} f_{1\,e}(\delta_e(\mathbf{x})) = 1 \cdot 1 + 1 \cdot 15 + 1 \cdot 1 = 17 < 18 = 9 \cdot 2 = \sum_{e \in P_6} f_{1\,e}(\delta_e(\mathbf{x})).$$

Furthermore, player 2 is satisfied since

$$\begin{aligned} \sum_{e \in P_6} f_{2\,e}(\delta_e(\mathbf{x})) = 2 \cdot 2 = 4 &= 4 = 4 \cdot 1 + 1 \cdot 0 = \sum_{e \in P_2} f_{2\,e}(\delta_e(\mathbf{x})), \\ \sum_{e \in P_6} f_{2\,e}(\delta_e(\mathbf{x})) = 2 \cdot 2 = 4 &= 4 = 1 \cdot 0 + 4 \cdot 1 = \sum_{e \in P_4} f_{2\,e}(\delta_e(\mathbf{x})). \end{aligned}$$

Obviously, players 3 and 4 do not want to deviate and thus $\mathbf{x}$ is a Wardrop equilibrium.

- Consider now the other strategy profile $\mathbf{y}$. Player 1 is satisfied since

$$\sum_{e \in P_3} f_{1\,e}(\delta_e(\mathbf{y})) = 1 \cdot \frac{3}{4} + 1 \cdot \frac{57}{4} + 1 \cdot \frac{3}{4} = \frac{63}{4} = 9 \cdot \frac{7}{4} = \sum_{e \in P_6} f_{1\,e}(\delta_e(\mathbf{y})).$$

Player 2 does not want to deviate since

$$\begin{aligned} \sum_{e \in P_2} f_{2\,e}(\delta_e(\mathbf{y})) = 4 \cdot \frac{3}{4} + \frac{1}{2} = \frac{7}{2} &= \frac{7}{2} = 1 \cdot \frac{1}{2} + 4 \cdot \frac{3}{4} = \sum_{e \in P_4} f_{2\,e}(\delta_e(\mathbf{y})), \\ \sum_{e \in P_2} f_{2\,e}(\delta_e(\mathbf{y})) = 4 \cdot \frac{3}{4} + \frac{1}{2} = \frac{7}{2} &= \frac{7}{2} = 2 \cdot \frac{7}{4} = \sum_{e \in P_6} f_{2\,e}(\delta_e(\mathbf{y})). \end{aligned}$$

It is again obvious that players 3 and 4 are also satisfied and thus $\mathbf{y}$ is a Wardrop equilibrium.

Having established that $\mathbf{x}$ and $\mathbf{y}$ are Wardrop equilibria we can now prove both parts (a) and (b) of the theorem.

Part (a): It is easy to see that $\delta_e(\mathbf{x}) \neq \delta_e(\mathbf{y})$ for all edges $e \in E \setminus \{e_7, e_8\}$. For the social cost of the Wardrop equilibrium $\mathbf{x}$ observe that

$$\begin{aligned} x_{1P_3} \sum_{e \in P_3} a_{1e} \delta_e(\mathbf{x}) &= 1 \cdot (1 + 15 + 1) = 17, \\ x_{2P_6} \sum_{e \in P_6} a_{2e} \delta_e(\mathbf{x}) &= 1 \cdot 2 \cdot 2 = 4, \\ x_{3P_6} \sum_{e \in P_6} a_{3e} \delta_e(\mathbf{x}) &= 1 \cdot 1 \cdot 2 = 2, \\ x_{4P_7} \sum_{e \in P_7} a_{4e} \delta_e(\mathbf{x}) &= 14 \cdot (14 + 15 + 14) = 602, \end{aligned}$$

and hence $\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon) = 17 + 4 + 2 + 602 = 625$ whereas we have for the social cost of the Wardrop equilibrium $\mathbf{y}$ that

$$\begin{aligned}
y_{1P_3} \sum_{e \in P_3} a_{1e}\delta_e(\mathbf{y}) + y_{1P_6} \sum_{e \in P_6} a_{1e}\delta_e(\mathbf{y}) &= \frac{1}{4} \cdot 1 \cdot \frac{63}{4} + \frac{3}{4} \cdot 9 \cdot \frac{7}{4} = \frac{63}{4}, \\
y_{2P_2} \sum_{e \in P_2} a_{2e}\delta_e(\mathbf{y}) + y_{2P_4} \sum_{e \in P_4} a_{2e}\delta_e(\mathbf{y}) &= \frac{1}{2}\left(3 + \frac{1}{2}\right) + \frac{1}{2}\left(\frac{1}{2} + 3\right) = \frac{7}{2}, \\
y_{3P_6} \sum_{e \in P_6} a_{3e}\delta_e(\mathbf{y}) &= 1 \cdot 1 \cdot \frac{7}{4} = \frac{7}{4}, \\
y_{4P_7} \sum_{e \in P_7} a_{4e}\delta_e(\mathbf{y}) &= 14 \cdot \left(14 + \frac{57}{4} + 14\right) = \frac{1183}{2},
\end{aligned}$$

and thus $\mathsf{SC}_{\mathsf{TL}}(\mathbf{y}, \Upsilon) = \frac{63}{4} + \frac{7}{2} + \frac{7}{4} + \frac{1183}{2} = \frac{1225}{2} \neq 625 = \mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon)$.

Part (b): Now consider the strategy profile $\mathbf{z} = \frac{1}{3} \cdot \mathbf{x} + \frac{2}{3} \cdot \mathbf{y}$ where $z_{1P_3} = z_{1P_6} = \frac{1}{2}$, $z_{2P_2} = z_{2P_4} = z_{2P_6} = \frac{1}{3}$, $z_{3P_6} = 1$, and $z_{4P_7} = 14$. Observe that $\mathbf{z}$ is no Wardrop equilibrium since the player-specific latency of player 1 on the path P_3 is $1 \cdot \frac{5}{6} + 1 \cdot \frac{29}{2} + 1 \cdot \frac{5}{6} = \frac{97}{6}$ whereas it is $9 \cdot \frac{11}{6} = \frac{33}{2}$ on P_6. This shows that the set of Wardrop equilibria of this instance does not form a convex set. ■

Theorem 7.5: *Consider a Wardrop game Υ with strictly increasing player-specific latency functions on parallel links and two associated Wardrop equilibria $\mathbf{x}$, $\mathbf{y}$. Then,*

(a) $\delta_e(\mathbf{x}) = \delta_e(\mathbf{y})$ for all $e \in E$ and $\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon) = \mathsf{SC}_{\mathsf{TL}}(\mathbf{y}, \Upsilon)$,

(b) the set of Wardrop equilibria for Υ forms a convex set.

Proof: We first show (a). Afterwards we will use (a) to show (b).

Part (a): We first of all show that $\delta_e(\mathbf{x}) = \delta_e(\mathbf{y})$ for all $e \in E$. Assume, by way of contradiction, that $\delta_e(\mathbf{x}) \neq \delta_e(\mathbf{y})$ for some $e \in E$ and let $M = \{e \in E;\ \delta_e(\mathbf{y}) < \delta_e(\mathbf{x})\}$. Fix two edges $r \notin M$, $j \in M$ and let $i \in [n]$ be some player with $x_{ij} > 0$. Then,

$$f_{ij}(\delta_j(\mathbf{y})) \ < \ f_{ij}(\delta_j(\mathbf{x})) \ \leq \ f_{ir}(\delta_r(\mathbf{x})) \ \leq \ f_{ir}(\delta_r(\mathbf{y})), \tag{7.1}$$

where one has to recall for the first inequality that $j \in M$ and that f_{ij} is strictly increasing, for the second inequality that $x_{ij} > 0$ and $\mathbf{x}$ is a Wardrop equilibrium, and for the third inequality that $r \notin M$. Since $\mathbf{y}$ is a Wardrop equilibrium we get with (7.1) that $y_{ir} = 0$. It follows that every player i who assigns in $\mathbf{x}$ non-zero traffic to at least one edge in

$j \in M$ assigns in $\mathbf{y}$ all its traffic to edges in M. This is a contradiction since we assumed that an edge $e \in E$ with $\delta_e(\mathbf{x}) \neq \delta_e(\mathbf{y})$ and thus also an edge $e \in M$ with $\delta_e(\mathbf{y}) < \delta_e(\mathbf{x})$ exists.

For the proof of $\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon) = \mathsf{SC}_{\mathsf{TL}}(\mathbf{y}, \Upsilon)$ let $\delta_e = \delta_e(\mathbf{x}) = \delta_e(\mathbf{y})$ denote the load on $e \in E$ in all Wardrop equilibria. Consider now a player $i \in [n]$. In both Wardrop equilibria $\mathbf{x}$ and $\mathbf{y}$ the player i only uses edges $k \in E$ for which $f_{ik}(\delta_k) = \min_{e \in E}\{f_{ie}(\delta_e)\}$. Therefore the social cost in both Wardrop equilibria $\mathbf{x}$ and $\mathbf{y}$ is given by:

$$\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon) \;=\; \mathsf{SC}_{\mathsf{TL}}(\mathbf{y}, \Upsilon) \;=\; \sum_{i \in [n]} w_i \cdot \min_{e \in E}\{f_{ie}(\delta_e)\}.$$

Part (b): Let again $\delta_e = \delta_e(\mathbf{x}) = \delta_e(\mathbf{y})$ denote the load on $e \in E$ in all Wardrop equilibria. Since $\mathbf{x}$ and $\mathbf{y}$ are Wardrop equilibria we have:

$$\begin{aligned} x_{ie} > 0 \;&\Rightarrow\; f_{ie}(\delta_e) \leq f_{ik}(\delta_k)\;\; \forall\, k \in E \quad \text{and} \\ y_{ie} > 0 \;&\Rightarrow\; f_{ie}(\delta_e) \leq f_{ik}(\delta_k)\;\; \forall\, k \in E. \end{aligned}$$

Putting these two conditions together we get:

$$(x_{ie} > 0) \vee (y_{ie} > 0) \quad \Rightarrow \quad f_{ie}(\delta_e) \leq f_{ik}(\delta_k)\;\; \forall\, k \in E. \tag{7.2}$$

Define now the strategy profile $\mathbf{z}$ by setting $z_{ie} = \lambda \cdot x_{ie} + (1 - \lambda) \cdot y_{ie}$ for some λ, $0 < \lambda < 1$, and all $i \in [n], e \in E$. Then we have for all $e \in E$ that:

$$\delta_e(\mathbf{z}) \;=\; \sum_{i=1}^{n} z_{ie} \;=\; \lambda \cdot \delta_e + (1 - \lambda) \cdot \delta_e \;=\; \delta_e.$$

Furthermore, $z_{ie} > 0$ if and only if $(x_{ie} > 0) \vee (y_{ie} > 0)$. Thus condition (7.2) shows that $\mathbf{z}$ is a Wardrop equilibrium. ■

7.4 Price of Anarchy

In this section we give bounds on the price of anarchy. The proof of the upper bound uses the same technique as the proof of Theorem 6.4.

Theorem 7.6 (Gairing, Monien, Tiemann [47]): *Consider a Wardrop game Υ with player-specific affine latency functions and an associated Wardrop equilibrium $\mathbf{x}$. Then,*

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon)}{\mathsf{OPT}_{\mathsf{TL}}(\Upsilon)} \;\leq\; \begin{cases} \frac{4}{4 - \Delta(\Upsilon)} & \text{if } \Delta(\Upsilon) \leq 2, \\ \Delta(\Upsilon) & \text{otherwise.} \end{cases}$$

Theorem 7.7 (Gairing, Monien, Tiemann [47]): *For each $n \in \mathbb{N}$ there is a Wardrop game Υ with n players, player-specific linear latency functions, and $\Delta(\Upsilon) = n^2$ on 2 parallel links that possesses a Wardrop equilibrium $\mathbf{x}$ such that*

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\mathbf{x}, \Upsilon)}{\mathsf{OPT}_{\mathsf{TL}}(\Upsilon)} \geq \frac{1}{4} \cdot \sqrt{\Delta(\Upsilon)}.$$

For Wardrop games with affine latency functions, Roughgarden and Tardos [85] showed that the price of anarchy is exactly $\frac{4}{3}$. Theorem 7.6 with $\Delta(\Upsilon) = 1$ implies that the price of anarchy does not change even if the affine latency functions of the players have player-specific constants $b_{ie} \neq b_{ke}$. Although our upper bound is tight for $\Delta(\Upsilon) = 1$ there is for large $\Delta(\Upsilon)$ still a gap between the upper bound of $\Delta(\Upsilon)$ and the lower bound.

7.5 Existence of Wardrop Equilibria

Every Wardrop game possesses a Wardrop equilibrium (see [12]). It is possible to use Brouwer's fixed point theorem to prove the existence of equilibria for our more general class of games.

Theorem 7.8 (Gairing, Monien, Tiemann [47]): *Every Wardrop game Υ with strictly increasing player-specific latency functions possesses a Wardrop equilibrium.*

7.6 Conclusion and Directions for Further Research

In this chapter, we considered Wardrop games with player-specific affine latency functions. In this setting, we were able to derive results on the convergence to equilibria, the polynomial time computation of equilibria, and the price of anarchy. However, several challenging problems remain open. We conclude this chapter by stating the most important ones:

- Based on a minimization of our new potential function Ψ a Wardrop equilibrium can be computed in polynomial time for the case of player-specific linear latency functions and parallel links (Theorem 7.3). Although Theorem 7.4 shows that a similar argumentation is impossible if we do not require parallel links we do not know whether there is a comparable function for a setting with more general (e.g. affine) latency functions.

- For the price of anarchy, we showed an upper bound of $\Delta(\Upsilon)$ if $\Delta(\Upsilon) > 2$ (Theorem 7.6) and a lower bound of $\frac{1}{4} \cdot \sqrt{\Delta(\Upsilon)}$ (Theorem 7.7). Hence the exact price of anarchy is still an open problem.

CHAPTER 8

Bayesian Routing Games

8.1 Introduction

In many large-scale, non-cooperative systems, users have only incomplete information about the system for several reasons. In his honored work, Harsanyi [51] introduced an elegant approach that can be used to study non-cooperative games with *incomplete information*, where the players are uncertain about some parameters. The *Harsanyi transformation* converts such a game with incomplete information to a game where players have different *types*. The type of a player represents its private information that is not common knowledge to all players.

In the resulting *Bayesian game*, each player's uncertainty about each other's type is described by a probability distribution over all possible *type profiles*. Using this probability distribution, players make their decisions according to Bayesian decision theory [15]. In Bayesian decision theory, probabilities are used to measure the degree of belief that a person has in some proposition.

In this chapter, we follow Harsanyi's approach and study the *Bayesian routing games* and *weighted Bayesian congestion games* that we introduced in Section 3.3. We allow here that the players do not know each other's weight. Thus there is for each player a set of possible types and each type of a player corresponds to some weight. If each player of a *Bayesian routing game* (or a *weighted Bayesian congestion game*) has only a single type, so that players are completely informed about each other's weight, then we are in the setting of weighted congestion games on parallel links (or weighted congestion games) like they were introduced in Section 3.1.

8.1.1 Contribution

Due to the new dimension that the incomplete information introduces to routing games, the analysis of the Bayesian routing game requires new techniques. In this chapter, we introduce such techniques and we present a comprehensive collection of results for the Bayesian routing game. We partition our results into three major parts:

(1) Existence and computational complexity of pure Bayesian Nash equilibria:

Our equilibria existence result applies for the class of *weighted Bayesian congestion games*. We define a new potential function that we use to prove that every weighted Bayesian congestion game possesses a pure Bayesian Nash equilibrium.

For the case of Bayesian routing games, identical links, and independent type distribution, we show that a pure Bayesian Nash equilibrium can be computed in polynomial time. This computation is based on Graham's LPT scheduling algorithm [50]. For the case of related links and independent type distribution, and also for the case of identical links and arbitrary type distribution, the complexity of computing a pure Bayesian Nash equilibrium remains open.

(2) Properties of fully mixed Bayesian Nash equilibria:

We show that for Bayesian routing games on identical links, the private cost of each player is maximized in a fully mixed Bayesian Nash equilibrium. This also implies that a player has the same private cost in any fully mixed Bayesian Nash equilibrium.

We define a certain fully mixed Bayesian Nash equilibrium that always exists. We show that, in general, there can exist more than one fully mixed Bayesian Nash equilibrium, and we study their structural properties. Finally, we determine the dimension of the space of fully mixed Bayesian Nash equilibria for the case of independent type distributions.

(3) Bounds on the price of anarchy:

We conclude with bounds on the price of anarchy for Bayesian routing games on identical links and three different social cost measures.

- For the expected maximum latency on a link as social cost measure, we show lower and upper bounds on the price of anarchy for different special cases. The exact value of the price of anarchy for this social cost measure remains open.
- A social cost measure that describes average player welfare is the sum of private costs. In this setting, it follows that each fully mixed Bayesian Nash equilibrium has maximum social cost. Using this

fact, we prove an upper bound of $\frac{m+n-1}{m}$ on the price of anarchy. We prove that this bound is asymptotically tight, already for complete information routing games.

– We also study social cost as maximum private cost. We show asymptotically tight upper bounds on the price of anarchy of $\frac{m+n-1}{m}$ for Bayesian routing games and of $2 - \frac{1}{m}$ for complete information routing games.

To the best of our knowledge, this is the first time that mixed Bayesian Nash equilibria are studied in combination with social cost.

8.1.2 Related Work

The related work that is relevant for this chapter focuses on congestion games on parallel links (Section 4.2), the existence of Nash equilibria in congestion games (Section 4.1), and Harsanyi's Bayesian games (Section 4.7).

8.1.3 Road Map

The rest of this chapter is organized as follows. Pure Bayesian Nash equilibria are studied in Section 8.2. Fully mixed Bayesian Nash equilibria are treated in Section 8.3. Section 8.4 studies the price of anarchy. We conclude in Section 8.5.

8.2 Pure Bayesian Nash Equilibria

In this section, we study the existence and the computational complexity of pure Bayesian Nash equilibria. We first show that Bayesian routing games and even weighted Bayesian congestion games are guaranteed to possess a pure Bayesian Nash equilibrium (Theorem 8.1). Furthermore, we give a polynomial time algorithm that is able to compute a pure Bayesian Nash equilibrium for each Bayesian routing game with identical links and independent type distribution (Theorem 8.2). Finally, we show that this algorithm does not work if a setting with related links (Proposition 8.1) or correlated type distribution (Proposition 8.2) is considered.

8.2.1 Existence of Pure Bayesian Nash Equilibria

To show that every weighted Bayesian congestion game with affine latency functions possesses a pure Bayesian Nash equilibrium we introduce a potential function Φ. For a pure strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ this function Φ is given by

$$\Phi(\sigma) = \sum_{i\in[n]} \sum_{t\in T_i} \sum_{e\in\sigma_i(t)} p(i,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-i}(\sigma,(\mathbf{p}|t_i=t)) + w(t)) + g_e(w(t))\right].$$

We consider in [46] an unilateral strategy change of a type agent that decreases the type agent's private cost and show that the potential function Φ also decreases. Thus we get that:

Theorem 8.1 (Gairing, Monien, Tiemann [46]): *Every weighted Bayesian congestion game with affine latency functions has a pure Bayesian Nash equilibrium.*

This generalizes a result of Fotakis et al. [39, Theorem 1] to the Bayesian setting. In particular our function Φ reduces to their potential function if each player has only a single type.

8.2.2 Computation of Pure Bayesian Nash Equilibria

We now turn to the model of Bayesian routing games on identical links and show how a pure Bayesian Nash equilibrium can be computed in polynomial time if the type distribution is independent. Given a Bayesian routing game Γ our algorithm constructs a complete information routing game Γ_{CI} (i.e., a weighted congestion game on parallel links), computes a Nash equilibrium α for this game Γ_{CI}, and uses α to obtain the Bayesian Nash equilibrium for Γ:

Algorithm 1 Pure Bayesian Nash equilibrium computation

Input: Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with $\mathbf{p}$ independent
Output: Pure Bayesian Nash equilibrium $\sigma = (\sigma_1, \dots, \sigma_n)$ for Γ

1: Calculate for each player $i \in [n]$ its expected weight $W(i)$.
2: Set $w(t'_i) = W(i)$ for all $i \in [n]$ and construct a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t'_1, ..., t'_n)\}, 1)$.
3: Compute a pure Nash equilibrium $\alpha : [n] \to [m]$ for Γ_{CI} in polynomial time with the LPT scheduling algorithm which assigns the players in order of non-increasing player weights to minimum load links (see [38, 50]).
4: Set $\sigma_i(t) = \alpha(i)$ for all players $i \in [n]$ and types $t \in T_i$.

A simple proof by contradiction shows that this algorithm indeed computes a pure Bayesian Nash equilibrium. Hence we get:

Theorem 8.2 (Gairing, Monien, Tiemann [46]): *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game on identical links with independent type distribution. Then, a normal pure Bayesian Nash equilibrium for Γ can be computed in time polynomial in the size of Γ even if $\mathbf{p}$ is represented in a compact form by a set of probabilities $p(i,t)$ for $i \in [n]$ and $t \in T_i$.*

The just given algorithm cannot be used to compute pure Bayesian Nash equilibria for the more general classes of Bayesian routing games either on related links or with correlated type distribution. The reason is that the algorithm always computes a *normal* Bayesian Nash equilibrium, whereas the following counter-examples show that a normal Bayesian Nash equilibrium does not exist in general.

Proposition 8.1: *There is a Bayesian routing game* Γ *on related links with independent type distribution that does not have a normal pure Bayesian Nash equilibrium.*

Proof: Consider the Bayesian routing game $\Gamma = (2, 2, \mathbf{c}, T_1 \times T_2, \mathbf{p})$ with two links of capacity $c_1 = 1$ and $c_2 = 5$. The two players have type sets $T_1 = \{t_1, t_1'\}$ and $T_2 = \{t_2\}$, where $w(t_1) = 1$, $w(t_1') = 5$, $w(t_2) = 10$, and $p(1, t_1) = p(1, t_1') = \frac{1}{2}$. We will now study the structure of pure Bayesian Nash equilibria for Γ and finally recognize that Γ has no *normal* pure Bayesian Nash equilibrium.

Let σ be an arbitrary pure Bayesian Nash equilibrium. Then,

$$\lambda^1_{(2,t_2)}(\sigma, \mathbf{p}) = \frac{\delta_1^{-2}(\sigma, \mathbf{p}) + w(t_2)}{c_1} \geq \frac{w(t_2)}{c_1} = 10$$

while

$$\lambda^2_{(2,t_2)}(\sigma, \mathbf{p}) = \frac{\delta_2^{-2}(\sigma, \mathbf{p}) + w(t_2)}{c_2} \leq \frac{\frac{1}{2} \cdot w(t_1) + \frac{1}{2} \cdot w(t_1') + w(t_2)}{c_2} = \frac{13}{5} < 10.$$

Thus, σ assigns t_2 to link 2, so $\sigma_2(t_2) = 2$. Consider now the types of player 1. We have

$$\lambda^1_{(1,t_1)}(\sigma, \mathbf{p}) = \frac{w(t_1)}{c_1} = 1 \quad \text{and} \quad \lambda^2_{(1,t_1)}(\sigma, \mathbf{p}) = \frac{w(t_2) + w(t_1)}{c_2} = \frac{11}{5},$$
$$\lambda^1_{(1,t_1')}(\sigma, \mathbf{p}) = \frac{w(t_1')}{c_1} = 5 \quad \text{and} \quad \lambda^2_{(1,t_1')}(\sigma, \mathbf{p}) = \frac{w(t_2) + w(t_1')}{c_2} = 3.$$

So σ assigns t_1 to link 1 and t_1' to link 2. It follows that σ is the unique pure Bayesian Nash equilibrium. However, σ is not *normal*. The claim follows. ■

Proposition 8.2: *There is a Bayesian routing game* Γ *on identical links with correlated type distribution that does not have a normal pure Bayesian Nash equilibrium.*

Proof: Consider the Bayesian routing game $\Gamma = (3, 2, \mathbf{1}, T_1 \times T_2 \times T_3, \mathbf{p})$ with 2 identical links and 3 players where the type sets are $T_1 = \{t_1, t_1'\}$, $T_2 = \{t_2, t_2'\}$ and $T_3 = \{t_3, t_3'\}$. The types are of weight $w(t_1) = w(t_2') = w(t_3) = w(t_3') = 1$ and $w(t_1') = w(t_2) = 2$. The correlated distribution $\mathbf{p}$ is given by $p(t_1, t_2, t_3) = p(t_1', t_2', t_3') = \frac{1}{2}$.

Assume, by way of contradiction, that a *normal* pure Bayesian Nash equilibrium σ exists; so, $\sigma_1(t_1) = \sigma_1(t'_1)$, $\sigma_2(t_2) = \sigma_2(t'_2)$, and $\sigma_3(t_3) = \sigma_3(t'_3)$. Let $j \neq k$ be the two links. Without loss of generality, set $\sigma_1(t_1) = \sigma_1(t'_1) = j$. Then, clearly

$$\lambda^j_{(2,t'_2)}(\sigma, \mathbf{p}) \geq w(t'_1) + w(t'_2) = 3 \quad \text{while} \quad \lambda^k_{(2,t'_2)}(\sigma, \mathbf{p}) \leq w(t'_3) + w(t'_2) = 2.$$

Thus, $\sigma_2(t'_2) = k$; hence, $\sigma_2(t_2) = \sigma_2(t'_2) = k$ for all normal pure Bayesian Nash equilibria σ. For the types of player 3, note that

$$\lambda^j_{(3,t_3)}(\sigma, \mathbf{p}) = w(t_1) + w(t_3) = 2 \quad \text{while} \quad \lambda^k_{(3,t_3)}(\sigma, \mathbf{p}) = w(t_2) + w(t_3) = 3,$$
$$\lambda^j_{(3,t'_3)}(\sigma, \mathbf{p}) = w(t'_1) + w(t'_3) = 3 \quad \text{while} \quad \lambda^k_{(3,t'_3)}(\sigma, \mathbf{p}) = w(t'_2) + w(t'_3) = 2.$$

Since σ is a Bayesian Nash equilibrium, $\sigma_3(t_3) = j$ and $\sigma_3(t'_3) = k$. Hence, σ is not normal. A contradiction. ∎

8.3 Fully Mixed Bayesian Nash Equilibria

In this section, we study fully mixed Bayesian Nash equilibria for the case of identical links. We first prove a simple expression for the private cost of each player in a fully mixed Bayesian Nash equilibrium (Proposition 8.3). This result is useful to show that the private cost of each player is maximized in a fully mixed Bayesian Nash equilibrium (Proposition 8.4). Please note that this also implies that a player has the same private cost in any fully mixed Bayesian Nash equilibrium. We define a certain fully mixed Bayesian Nash equilibrium that always exists (Definition 8.1). Finally, we describe the structural properties of fully mixed Bayesian Nash equilibria (Proposition 8.5) and we determine the dimension of the space of fully mixed Bayesian Nash equilibria (Theorem 8.3) for the case of independent type distributions.

We start by proving a technical lemma that we will apply in the proof of Proposition 8.3.

Lemma 8.1: *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated mixed strategy profile $\mathbf{Q}$. Then, for each player $i \in [n]$,*

$$\sum_{j \in [m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \;\;=\;\; \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t).$$

Proof: Clearly,

$$\begin{aligned} &\sum_{j \in [m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \\ &\quad = \sum_{j \in [m]} \sum_{\sigma \in \Sigma} \prod_{s \in [n]} q(s, \sigma_s) \cdot \delta_j^{-i}(\sigma, (\mathbf{p}|t_i = t)) \end{aligned}$$

$$= \sum_{j\in[m]} \sum_{\sigma\in\Sigma} \prod_{s\in[n]} q(s,\sigma_s) \sum_{\substack{(t_1,\ldots,t_n)\in T: \\ t_i=t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t) \sum_{\substack{s\in[n]\setminus\{i\}: \\ \sigma_s(t_s)=j}} w(t_s)$$

$$= \sum_{\sigma\in\Sigma} \prod_{s\in[n]} q(s,\sigma_s) \cdot \sum_{\substack{(t_1,\ldots,t_n)\in T: \\ t_i=t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t) \cdot \sum_{s\in[n]\setminus\{i\}} w(t_s)$$

$$= \sum_{\substack{(t_1,\ldots,t_n)\in T: \\ t_i=t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t) \cdot \sum_{s\in[n]\setminus\{i\}} w(t_s)$$

$$= \sum_{s\in[n]\setminus\{i\}} \sum_{\substack{(t_1,\ldots,t_n)\in T: \\ t_i=t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t) \cdot w(t_s)$$

$$= \sum_{s\in[n]\setminus\{i\}} W(s|t_i=t). \qquad \blacksquare$$

We continue to prove a simple expression for the private cost of each player in a fully mixed Bayesian Nash equilibrium.

Proposition 8.3: *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Then for each player $i \in [n]$,*

$$\mathsf{PC}_i(\mathbf{F},\mathbf{p}) = \frac{W}{m} + \frac{m-1}{m} \cdot W(i).$$

Proof: Fix any player $i \in [n]$. Clearly, for any link $k \in \mathsf{support}_{F_i}(i) = [m]$, and by Lemma 8.1,

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{F},\mathbf{p}) &= \sum_{t\in T_i} p(i,t) \cdot v_{(i,t)}(\mathbf{F},\mathbf{p}) \\
&= \sum_{t\in T_i} p(i,t) \cdot \big(w(t) + \delta_k^{-i}(\mathbf{F},(\mathbf{p}|t_i=t))\big) \\
&= \sum_{t\in T_i} p(i,t) \cdot w(t) + \sum_{t\in T_i} p(i,t) \cdot \delta_k^{-i}(\mathbf{F},(\mathbf{p}|t_i=t)) \\
&= W(i) + \sum_{t\in T_i} p(i,t) \cdot \frac{1}{m} \cdot \sum_{j\in[m]} \delta_j^{-i}(\mathbf{F},(\mathbf{p}|t_i=t)) \\
&= W(i) + \frac{1}{m} \cdot \sum_{t\in T_i} p(i,t) \cdot \sum_{s\in[n]\setminus\{i\}} W(s|t_i=t) \qquad (8.1) \\
&= W(i) + \frac{1}{m} \cdot \sum_{s\in[n]\setminus\{i\}} \sum_{t\in T_i} p(i,t) \cdot W(s|t_i=t) \\
&= W(i) + \frac{1}{m} \cdot \sum_{s\in[n]\setminus\{i\}} W(s) \\
&= \frac{W}{m} + \frac{m-1}{m} \cdot W(i). \qquad \blacksquare
\end{aligned}$$

We now prove that the private cost of each player is maximized in a fully mixed Bayesian Nash equilibrium. For the special case of complete information routing games this result is already known [44].

Proposition 8.4: *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links, an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$*, and a Bayesian Nash equilibrium* $\mathbf{Q}$*. Then for each player* $i \in [n]$*,*

$$\mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) \quad \leq \quad \mathsf{PC}_i(\mathbf{F}, \mathbf{p}).$$

Proof: Fix any player $i \in [n]$. Then, for any link $j \in [m]$,

$$\begin{aligned} \mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) &= \sum_{t \in T_i} p(i,t) \cdot v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \\ &\leq \sum_{t \in T_i} p(i,t) \cdot \left(w(t) + \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \right), \end{aligned}$$

since $\mathbf{Q}$ is a Bayesian Nash equilibrium. In particular,

$$\begin{aligned} \mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) &\leq \sum_{t \in T_i} p(i,t) \cdot \left(w(t) + \min_{j \in [m]} \left\{ \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \right\} \right) \\ &\leq \sum_{t \in T_i} p(i,t) \cdot \left(w(t) + \frac{1}{m} \cdot \sum_{j \in [m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \right) \\ &= \sum_{t \in T_i} p(i,t) \cdot \left(w(t) + \frac{1}{m} \cdot \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t) \right) \\ &= W(i) + \frac{1}{m} \cdot \sum_{t \in T_i} p(i,t) \cdot \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t) \\ &= \mathsf{PC}_i(\mathbf{F}, \mathbf{p}), \end{aligned}$$

by Equation (8.1), as needed. ∎

We proceed to define a particular fully mixed strategy profile $\overline{\mathbf{F}}$.

Definition 8.1: *A* standard *fully mixed strategy profile* $\overline{\mathbf{F}}$ *is a fully mixed strategy profile that assigns every type agent to every link with probability* $\frac{1}{m}$*.*

It is easy to see that for any Bayesian routing game Γ on identical links, a standard fully mixed strategy profile is a Bayesian Nash equilibrium. For the special case of complete information routing games, this fact was also observed in [67].

In general, there exists more than one fully mixed Bayesian Nash equilibrium (see, e.g., Proposition 8.7). We now give an exact characterization of all fully mixed Bayesian Nash equilibria.

Proposition 8.5: *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links with independent type distribution and an associated fully mixed strategy profile* $\mathbf{F}$. *Then,* $\mathbf{F}$ *is a fully mixed Bayesian Nash equilibrium if and only if for all players* $i \in [n]$ *and links* $j \in [m]$ *it holds that*

$$\frac{1}{m} \cdot W(i) \;=\; \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i, t) \cdot w(t).$$

Proof: For any player $i \in [n]$ and link $j \in [m]$, set

$$\mu(\mathbf{F}, i, j) \;=\; \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i, t) \cdot w(t).$$

Observe that for any player $i \in [n]$ and link $j \in [m]$,

$$\begin{aligned}
&\delta_j^{-i}(\mathbf{F}, \mathbf{p}) \\
&= \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \cdot \sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) \cdot \sum_{\substack{k \in [n] \setminus \{i\}: \\ \sigma_k(t_k) = j}} w(t_k) \\
&= \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \cdot \sum_{k \in [n] \setminus \{i\}} \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \cdot \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma'_k \in \Sigma_k} f(k, \sigma'_k) \cdot \sum_{\substack{\sigma \in \Sigma: \\ \sigma_k = \sigma'_k}} \prod_{s \in [n] \setminus \{k\}} f(s, \sigma_s) \cdot \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma'_k \in \Sigma_k} f(k, \sigma'_k) \cdot \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j)\ .
\end{aligned}$$

Proof of $\Leftarrow$: Consider first an arbitrary fully mixed strategy profile $\mathbf{F}$ that satisfies for all players $i \in [n]$ and all links $j \in [m]$ that $\mu(\mathbf{F}, i, j) = \frac{1}{m} \cdot W(i)$. Then, for all players $i \in [n]$, types $t \in T_i$, and links $j \in [m]$,

$$\begin{aligned}
\lambda^j_{(i,t)}(\mathbf{F}, \mathbf{p}) &= \delta_j^{-i}(\mathbf{F}, (\mathbf{p} | t_i = t)) + w(t) \\
&= \delta_j^{-i}(\mathbf{F}, \mathbf{p}) + w(t) \\
&= \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j) + w(t) \\
&= \sum_{k \in [n] \setminus \{i\}} \frac{1}{m} \cdot W(k) + w(t).
\end{aligned}$$

Hence we get for the private cost of type agent (i, t),

$$\begin{aligned} v_{(i,t)}(\mathbf{F}, \mathbf{p}) &= \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \lambda_{(i,t)}^{\sigma_i(t)}(\mathbf{F}, \mathbf{p}) \\ &= \sum_{k \in [n] \setminus \{i\}} \frac{1}{m} \cdot W(k) + w(t). \end{aligned}$$

So, $v_{(i,t)}(\mathbf{F}, \mathbf{p}) = \lambda_{(i,t)}^{j}(\mathbf{F}, \mathbf{p})$ for all players $i \in [n]$, types $t \in T_i$, and links $j \in [m]$. Thus, $\mathbf{F}$ is a fully mixed Bayesian Nash equilibrium.

Proof of $\Rightarrow$: Assume now that $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium. Hence, $\mathsf{support}_{Q_i}(t) = [m]$ for all players $i \in [n]$ and types $t \in T_i$. Since $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium and $\mathbf{p}$ is independent, it follows that for all links $j \in \mathsf{support}_{Q_i}(t) = [m]$,

$$\begin{aligned} v_{(i,t)}(\mathbf{F}, \mathbf{p}) &= \lambda_{(i,t)}^{j}(\mathbf{F}, \mathbf{p}) \\ &= \delta_j^{-i}(\mathbf{F}, (\mathbf{p} | t_i = t)) + w(t) \\ &= \delta_j^{-i}(\mathbf{F}, \mathbf{p}) + w(t). \end{aligned}$$

So, for all players $i \in [n]$ and pair of links $j, l \in [m]$,

$$\delta_j^{-i}(\mathbf{F}, \mathbf{p}) = \delta_l^{-i}(\mathbf{F}, \mathbf{p}).$$

Since $\delta_j^{-i}(\mathbf{F}, \mathbf{p}) = \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j)$ for any player i and link j, it follows that for an arbitrary pair of players $i_1, i_2 \in [n]$ with $i_1 \neq i_2$ and an arbitrary pair of links $j_1, j_2 \in [m]$ with $j_1 \neq j_2$,

$$\sum_{k \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, k, j_1) = \sum_{k \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, k, j_2) \tag{8.2}$$

and

$$\sum_{k \in [n] \setminus \{i_2\}} \mu(\mathbf{F}, k, j_1) = \sum_{k \in [n] \setminus \{i_2\}} \mu(\mathbf{F}, k, j_2). \tag{8.3}$$

Subtracting (8.3) from (8.2) yields that

$$\mu(\mathbf{F}, i_2, j_1) - \mu(\mathbf{F}, i_1, j_1) = \mu(\mathbf{F}, i_2, j_2) - \mu(\mathbf{F}, i_1, j_2),$$

or equivalently

$$0 = \mu(\mathbf{F}, i_2, j_1) - \mu(\mathbf{F}, i_2, j_2) + \mu(\mathbf{F}, i_1, j_2) - \mu(\mathbf{F}, i_1, j_1).$$

Summing up over all players $i_2 \in [n] \setminus \{i_1\}$ yields that

$$\begin{aligned}
0 &= \sum_{i_2 \in [n]\setminus\{i_1\}} \mu(\mathbf{F}, i_2, j_1) - \sum_{i_2 \in [n]\setminus\{i_1\}} \mu(\mathbf{F}, i_2, j_2) \\
&\quad + \sum_{i_2 \in [n]\setminus\{i_1\}} \mu(\mathbf{F}, i_1, j_2) - \sum_{i_2 \in [n]\setminus\{i_1\}} \mu(\mathbf{F}, i_1, j_1) \\
&= \delta_{j_1}^{-i_1}(\mathbf{F}, \mathbf{p}) - \delta_{j_2}^{-i_1}(\mathbf{F}, \mathbf{p}) \\
&\quad + (n-1) \cdot \mu(\mathbf{F}, i_1, j_2) - (n-1) \cdot \mu(\mathbf{F}, i_1, j_1) \\
&= (n-1) \cdot (\mu(\mathbf{F}, i_1, j_2) - \mu(\mathbf{F}, i_1, j_1)) .
\end{aligned}$$

It follows that for all players $i_1 \in [n]$ and pair of links $j_1, j_2 \in [m]$,

$$\mu(\mathbf{F}, i_1, j_1) = \mu(\mathbf{F}, i_1, j_2).$$

Clearly, for any player $i \in [n]$,

$$\begin{aligned}
W(i) &= \sum_{t \in T_i} p(i,t) \cdot w(t) \\
&= \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{t \in T_i} p(i,t) \cdot w(t) \\
&= \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{j \in [m]} \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i,t) \cdot w(t) \\
&= \sum_{j \in [m]} \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i,t) \cdot w(t) \\
&= \sum_{j \in [m]} \mu(\mathbf{F}, i, j) \\
&= m \cdot \mu(\mathbf{F}, i, j),
\end{aligned}$$

for any link $j \in [m]$. This implies that for all players $i \in [n]$ and links $j \in [m]$,

$$\mu(\mathbf{F}, i, j) = \frac{1}{m} \cdot W(i),$$

or

$$\frac{1}{m} \cdot W(i) = \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i,t) \cdot w(t),$$

as needed. ■

We finally determine a lower bound on the dimension of the space of fully mixed Bayesian Nash equilibria.

Theorem 8.3 (Gairing, Monien, Tiemann [46]): *Consider a Bayesian routing game Γ on identical links with independent type distribution. Then, the dimension of the space of fully mixed Bayesian Nash equilibria for Γ is at least $\sum_{i \in [n]} m^{|T_i|} - nm$.*

8.4 Social Cost and Price of Anarchy

In this section, we present bounds on the price of anarchy for three different social cost measures and Bayesian routing games on *identical* links. We start with social cost as expected maximum latency (Section 8.4.1), proceed with social cost as sum of private costs (Section 8.4.2), and close with social cost as maximum of private costs (Section 8.4.2).

8.4.1 Social Cost as Expected Maximum Latency

We will now study social cost as the expected maximum latency and show lower and upper bounds on the price of anarchy for different special cases.

For complete information routing games (i.e., weighted congestion games on parallel links) social cost as expected maximum latency was introduced by Koutsoupias and Papadimitriou [60]. Asymptotic tight bounds on the price of anarchy were given by Czumaj and Vöcking [26] and Koutsoupias et al. [59]. Their techniques use Chernoff bounds to show that for identical links the quotient between the expected maximum load and the maximum expected load on a link is at most $\mathcal{O}(\frac{\log m}{\log \log m})$. We prove that similar techniques cannot be applied for Bayesian routing games to prove an upper bound on the price of anarchy which is better than $\mathcal{O}(m)$.

Proposition 8.6 (Gairing, Monien, Tiemann [46]): *For any $\epsilon > 0$, there is a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links with independent type distribution and an associated pure Bayesian Nash equilibrium σ with $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma)$, such that for each link $j \in [m]$,*

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma)}{\delta_j(\sigma, \mathbf{p})} \geq \frac{m}{1+\epsilon}.$$

We now turn our attention to standard fully mixed Bayesian Nash equilibria and prove:

Theorem 8.4: *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$. Then,*

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} \;=\; \mathcal{O}\left(\frac{\log m}{\log \log m}\right).$$

Proof: Consider an arbitrary type profile $t = (t_1, ..., t_n) \in T$. Given t, we define the game $\Gamma_{\mathsf{CI}}(t) = (n, m, \mathbf{1}, \{(t_1, ..., t_n)\}, 1)$. Recall that for this complete information routing game $\Gamma_{\mathsf{CI}}(t)$, the unique fully mixed Nash equilibrium $\overline{\mathbf{Q}}(t)$ assigns each player to each link with probability $1/m$ (see [67, Lemma 15]). By [59, Theorem 4.4] or [26, Theorem 1.1], it holds that

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{Q}}, \Gamma_{\mathsf{CI}}(t))}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}(t))} = \mathcal{O}\left(\frac{\log m}{\log \log m}\right).$$

Recall that $\overline{\mathbf{F}}$ assigns every type agent to every link with probability $\frac{1}{m}$. Thus,

$$\begin{aligned}
\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma) &= \sum_{t \in T} p(t) \cdot \sum_{(\sigma_1(t_1), \ldots, \sigma_n(t_n)) \in [m]^n} \left(\frac{1}{m}\right)^n \cdot \max_{j \in [m]} \left\{ \sum_{\substack{i \in [n]: \\ \sigma_i(t_i) = j}} w(t_i) \right\} \\
&= \sum_{t \in T} p(t) \cdot \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{Q}}(t), \Gamma_{\mathsf{CI}}(t)) \\
&= \sum_{t \in T} p(t) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}(t)) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right) \\
&= \mathsf{OPT}_{\mathsf{MSP}}(\Gamma) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right),
\end{aligned}$$

as needed. ■

Theorem 8.4 implies that for standard fully mixed Nash equilibria, incomplete information has no impact on the price of anarchy when social cost is taken as expected maximum latency.

Since, in general, there is more than one fully mixed Bayesian Nash equilibrium, the natural question arises whether they have all the same expected maximum latency. As we see now, this is not the case.

Proposition 8.7: *There exists a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$ such that*

$$\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) \;>\; \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma).$$

Proof: Consider the Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with $n = 2, m = 3$ and $T_i = \{t_i, t'_i\}$ with $w(t_i) = 2, w(t'_i) = 1$ for all players $i \in \{1, 2\}$; set $p(i, t_i) = p(i, t'_i) = \frac{1}{2}$ for all players $i \in \{1, 2\}$. Consider the standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$ and some other fully mixed Bayesian Nash equilibrium $\mathbf{F}$ which we define below:

- $\overline{\mathbf{F}}$ assigns each type to each link with a probability of $\frac{1}{3}$. Thus, the two players are assigned to the same link with a probability of $\frac{1}{3}$. In this case, the maximum latency can be 2, 3, or 4. With a probability of $\frac{2}{3}$, the players are assigned to different links. In this case the maximum latency can be 1 or 2. Hence, the social cost of the standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$ is

$$\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma) = \frac{1}{3} \cdot \left(\frac{1}{4} \cdot 2 + \frac{1}{2} \cdot 3 + \frac{1}{4} \cdot 4\right) + \frac{2}{3} \cdot \left(\frac{1}{4} \cdot 1 + \frac{3}{4} \cdot 2\right) = \frac{13}{6}.$$

- The fully mixed strategy profile $\mathbf{F}$ assigns each type of weight 1 to link 1 with a probability of $\frac{1}{2}$, to link 2 with a probability of $\frac{1}{4}$ and to link

3 with a probability of $\frac{1}{4}$. Each type of weight 2 is assigned to link 1 with a probability of $\frac{1}{4}$, to link 2 with a probability of $\frac{3}{8}$ and to link 3 with a probability of $\frac{3}{8}$. Observe that for all $i \in \{1,2\}$ we get $\delta_1^{-i}(\mathbf{F}) = \frac{1}{2} \cdot \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{4} \cdot 2 = \frac{1}{2}$ and $\delta_2^{-i}(\mathbf{F}) = \delta_3^{-i}(\mathbf{F}) = \frac{1}{2} \cdot \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot \frac{3}{8} \cdot 2 = \frac{1}{2}$. Thus, $\mathbf{F}$ is a Bayesian Nash equilibrium.

With probability $\frac{1}{4}$ both players are of weight 1. In this case they use the same link with probability $(\frac{1}{2})^2 + 2 \cdot (\frac{1}{4})^2 = \frac{3}{8}$. With probability $\frac{1}{2}$, exactly one of the two players is of weight 1. In this case, the players use the same link with probability $\frac{1}{2} \cdot \frac{1}{4} + 2 \cdot \frac{1}{4} \cdot \frac{3}{8} = \frac{5}{16}$. With the remaining probability $\frac{1}{4}$, both players are of weight 2. In this case, the players use the same link with probability $(\frac{1}{4})^2 + 2 \cdot (\frac{3}{8})^2 = \frac{11}{32}$. Hence we get that the social cost of $\mathbf{F}$ is

$$
\begin{aligned}
&\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) \\
&= \frac{1}{4} \cdot \left(\frac{3}{8} \cdot 2 + \frac{5}{8} \cdot 1\right) + \frac{1}{2} \cdot \left(\frac{5}{16} \cdot 3 + \frac{11}{16} \cdot 2\right) + \frac{1}{4} \cdot \left(\frac{11}{32} \cdot 4 + \frac{21}{32} \cdot 2\right) \\
&= \frac{139}{64}.
\end{aligned}
$$

Observe that $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) = \frac{139}{64} = \frac{417}{192} > \frac{416}{192} = \frac{13}{6} = \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma)$. ■

It is known (see [65, Section 8.E]) that mixed Nash equilibria in games with complete information are related to pure Bayesian Nash equilibria in a Bayesian game, where for each player all its types are identical. The following definition and theorem applies this to Bayesian routing games.

Definition 8.2: *A* CI-like game *is a Bayesian routing game with an independent type distribution such that $w(t) = w(t')$ for all types $t, t' \in T_i$, where $i \in [n]$.*

We call these games CI-like games (where CI stands for complete information) since they are similar to complete information routing games in the sense that the weight of a player does not depend on its type. For complete information routing games, there exist asymptotically tight upper bounds on the price of anarchy for the cases of identical links [26, 59] and related links [26]. We use these bounds to prove:

Theorem 8.5: *Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ be a CI-like game with an associated pure Bayesian Nash equilibrium σ. Then*

(a) $\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log\log m}\right)$*, for the case of identical links,*

(b) $\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log\log\log m}\right)$*, for the case of related links,*

and there are CI-like games for which both bounds are asymptotically tight.

Proof: The proof is structured as follows: We first define a construction that maps any CI-like game Γ with an associated pure strategy profile σ to a complete information routing Γ_{CI} with associated (mixed) strategy profile $\mathbf{Q}$. For this construction, we show that $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma_{\mathsf{CI}})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$, and that $\mathbf{Q}$ is a Nash equilibrium if σ is a Bayesian Nash equilibrium. From these properties of our construction, we derive that the corresponding upper bounds on the price of anarchy [26, 59] for complete information routing games also hold for CI-like games.

To prove tightness, we show that for every complete information routing game Γ_{CI} with associated (mixed) Nash equilibrium $\mathbf{Q}$, we can define a CI-like game Γ with associated pure Bayesian Nash equilibrium σ, such that our construction maps Γ and σ to Γ_{CI} and $\mathbf{Q}$, respectively. This implies that also the lower bounds on the price of anarchy can be carried over to the CI-like games. We start by defining our construction.

Construction $\Gamma \mapsto \Gamma_{\mathsf{CI}}$: Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ be a CI-like game. For each $i \in [n]$, denote by $w_i = w(t)$ the weight of all types $t \in T_i$. Define a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T', 1)$ where $T' = \{(t'_1, \ldots, t'_n)\}$ and $w(t'_i) = w_i$ for all $i \in [n]$.

Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a pure strategy profile for the CI-like game Γ. Denote by Σ' the set of all pure strategy profiles for Γ_{CI}; thus, $\Sigma' = \Sigma'_1 \times \ldots \times \Sigma'_n$, where for each player $i \in [n]$, the set Σ'_i consists of all possible pure strategies $\sigma'_i : \{t'_i\} \to [m]$ for player i.

Define a mixed strategy profile $\mathbf{Q}$ for Γ_{CI}, where for each player $i \in [n]$ and all pure strategies $\sigma'_i \in \Sigma'_i$ the probability $q(i, \sigma'_i)$ is given by $q(i, \sigma'_i) = \sum_{t \in T_i : \sigma_i(t) = \sigma'_i(t'_i)} p(i, t)$. We proceed by showing properties of our construction.

- $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma_{\mathsf{CI}})$: To show that the strategy profiles σ for Γ and $\mathbf{Q}$ for Γ_{CI} are of the same social cost observe that

$$
\begin{aligned}
&\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) \\
&= \sum_{(t_1,\ldots,t_n)\in T} p(t_1, \ldots, t_n) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i) \right\} \\
&= \sum_{(t_1,\ldots,t_n)\in T} \prod_{i\in[n]} p(i, t_i) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w_i \right\} \\
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ \sigma_i(t_i)=\sigma'_i(t'_i)\forall i\in[n]}} \prod_{i\in[n]} p(i, t_i) \right) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i \right\}
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\prod_{i\in[n]} \sum_{\substack{t\in T_i:\\ \sigma_i(t)=\sigma'_i(t'_i)}} p(i,t) \right) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i \right\} \\
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \prod_{i\in[n]} q(i,\sigma'_i) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i \right\} \\
&= \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}}).
\end{aligned}
$$

- $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$: To show $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) \geq \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$ observe that our construction maps a pure strategy profile for Γ of optimum social cost to a strategy profile for Γ_{CI} that has the same social cost.

  For the other direction $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) \leq \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$, observe that there always exists a *pure* strategy profile $\hat{\sigma}'$ for Γ_{CI} of optimum social cost, i.e. $\mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}}) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$. Consider the normal pure strategy profile $\hat{\sigma}$ for Γ that assigns for each $i \in [n]$ all types of player i to the link to that $\hat{\sigma}'$ assigns player i, so $\hat{\sigma}_i(t) = \hat{\sigma}_i{}'(t'_i)$ for all players $i \in [n]$ and all types $t \in T_i$. Notice that our construction transforms Γ and $\hat{\sigma}$ back to Γ_{CI} and $\hat{\sigma}'$. Thus $\mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma},\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}})$. We get that

$$
\begin{aligned}
\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) &\leq \mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma},\Gamma) \\
&= \mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}}) \\
&= \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}).
\end{aligned}
$$

- Mapping of Equilibria: Clearly, for all players $i \in [n]$, types $t \in T_i$, and links $j \in [m]$,

$$
\begin{aligned}
&\lambda^j_{(i,t)}(\sigma,\mathbf{p}) \\
&= \frac{1}{c_j} \cdot \left(w(t) + \delta_j^{-i}(\sigma,\mathbf{p}) \right) \\
&= \frac{1}{c_j} \cdot \left(w(t) + \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w(t_s) \right) \\
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(t_1,\ldots,t_n)\in T} \prod_{s\in[n]} p(s,t_s) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w_s \right) \\
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ \sigma_s(t_s)=\sigma'_s(t'_s)\forall s\in[n]}} \prod_{s\in[n]} p(s,t_s) \right) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w_s \right)
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\prod_{s\in[n]} \sum_{\substack{t_s\in T_s:\\ \sigma_s(t_s)=\sigma'_s(t'_s)}} p(s,t_s) \right) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w_s \right) \\
&= \frac{1}{c_j} \cdot \left(w(t'_i) + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \prod_{s\in[n]} q(s,\sigma'_s) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w(t'_s) \right) \\
&= \frac{1}{c_j} \cdot \left(w(t'_i) + \delta_j^{-i}(\mathbf{Q},1) \right) \\
&= \lambda^j_{(i,t'_i)}(\mathbf{Q},1).
\end{aligned}
$$

We now use this property to show that $\mathbf{Q}$ is a Nash equilibrium for Γ_{CI} if σ is a pure Bayesian Nash equilibrium for Γ. So, let σ be a pure Bayesian Nash equilibrium for Γ. Fix an arbitrary player $i \in [n]$. Remember that in Γ all types of player i have the same weight. Thus,

$$
v_{(i,t)}(\sigma,\mathbf{p}) = v_{(i,\hat{t})}(\sigma,\mathbf{p})
$$

for all pairs of types $t, \hat{t} \in T_i$. Since σ is a pure Bayesian Nash equilibrium for Γ, this implies that for all types $t \in T_i$,

$$
\begin{aligned}
&v_{(i,t)}(\sigma,\mathbf{p}) = \lambda^j_{(i,t)}(\sigma,\mathbf{p}) \quad \text{for all } j \in \mathsf{support}_{\sigma_i}(i) \text{ and} \\
&v_{(i,t)}(\sigma,\mathbf{p}) \le \lambda^j_{(i,t)}(\sigma,\mathbf{p}) \quad \text{for all } j \notin \mathsf{support}_{\sigma_i}(i).
\end{aligned}
$$

By definition of $\mathbf{Q}$,

$$
\mathsf{support}_{\sigma_i}(i) = \mathsf{support}_{\mathbf{Q}_i}(t'_i).
$$

It follows that

$$
\begin{aligned}
&v_{(i,t'_i)}(\mathbf{Q},1) = \lambda^j_{(i,t'_i)}(\mathbf{Q},1) \quad \text{for all } j \in \mathsf{support}_{\mathbf{Q}_i}(t'_i) \text{ and} \\
&v_{(i,t'_i)}(\mathbf{Q},1) \le \lambda^j_{(i,t'_i)}(\mathbf{Q},1) \quad \text{for all } j \notin \mathsf{support}_{\mathbf{Q}_i}(t'_i),
\end{aligned}
$$

so that $\mathbf{Q}$ is a Nash equilibrium.

Upper bounds on price of anarchy: Recall that by our construction, we have that $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})$ and $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$. Thus, resorting to the corresponding upper bounds on the price of anarchy from [59] and [26], we get

$$
\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} = \begin{cases} \mathcal{O}\left(\frac{\log m}{\log\log m}\right), & \text{for identical links,} \\ \mathcal{O}\left(\frac{\log m}{\log\log\log m}\right), & \text{for related links.} \end{cases}
$$

This completes the proof of the upper bounds.

Tightness of the upper bounds: From [59] and [26], there exist complete information routing games Γ_{CI} with an associated mixed Nash equilibrium $\mathbf{Q}$ such that

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} = \begin{cases} \Omega\left(\frac{\log m}{\log\log m}\right) & \text{, for the case of identical links,} \\ \Omega\left(\frac{\log m}{\log\log\log m}\right) & \text{, for the case of related links.} \end{cases}$$

Let $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T', 1)$, $T' = \{(t'_1, \ldots, t'_n)\}$, be such a complete information routing game with an associated mixed Nash equilibrium $\mathbf{Q}$. With a slight abuse of notation, we denote $\mathbf{Q} = (q(i,j))_{i\in[n],j\in[m]}$ where $q(i,j)$ is the probability that type $t'_i \in T'_i$ is assigned to link $j \in [m]$.

We define a CI-like game $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ and an associated pure strategy profile σ as follows:

> For each player $i \in [n]$, T_i consists of $|\mathsf{support}_{\mathbf{Q}_i}(i)|$ types, where we have a type t_i^j for every link $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. For all players $i \in [n]$ and links $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$, define $p(i, t_i^j) = q(i,j)$ and $\sigma_i(t_i^j) = j$.

Notice that our construction $\Gamma \mapsto \Gamma_{\mathsf{CI}}$ transforms the CI-like game Γ with associated pure strategy profile σ back to the complete information routing game Γ_{CI} with associated (mixed) Nash equilibrium $\mathbf{Q}$. It follows that $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$ and $\lambda^l_{(i,t_i^j)}(\sigma,\mathbf{p}) = \lambda^l_{(i,t'_i)}(\mathbf{Q},1)$ for all players $i \in [n]$, for all links $l \in [m]$, and for all $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. Since $\mathbf{Q}$ is a Nash equilibrium we have

$$\begin{aligned} v_{(i,t'_i)}(\mathbf{Q},1) &= \lambda^j_{(i,t'_i)}(\mathbf{Q},1) \quad \text{for all } j \in \mathsf{support}_{\mathbf{Q}_i}(t'_i) \text{ and} \\ v_{(i,t'_i)}(\mathbf{Q},1) &\leq \lambda^j_{(i,t'_i)}(\mathbf{Q},1) \quad \text{for all } j \notin \mathsf{support}_{\mathbf{Q}_i}(t'_i). \end{aligned}$$

Furthermore, $\mathsf{support}_{\sigma_i}(i) = \mathsf{support}_{\mathbf{Q}_i}(i)$ for all $i \in [n]$, and $\lambda^l_{(i,t_i^j)}(\sigma,\mathbf{p}) = \lambda^l_{(i,t'_i)}(\mathbf{Q},1)$ for all players $i \in [n]$, for all links $l \in [m]$, and for all $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. It follows that σ is a pure Bayesian Nash equilibrium with

$$\begin{aligned} \frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} &= \frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} \\ &= \begin{cases} \Omega\left(\frac{\log m}{\log\log m}\right) & \text{, for the case of identical links,} \\ \Omega\left(\frac{\log m}{\log\log\log m}\right) & \text{, for the case of related links.} \end{cases} \end{aligned}$$

This completes the proof. ■

We conclude with a lower bound on the price of anarchy for normal pure Bayesian Nash equilibria.

Theorem 8.6: *There exists a sequence of Bayesian routing games Γ on identical links and associated normal pure Bayesian Nash equilibria σ such that*

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} \;=\; \Omega\left(\frac{\log m}{\log\log m}\right).$$

Proof: Let $m > 10^{15}$ be a perfect square, i.e., we have that $\sqrt{m} \in \mathbb{N}$. Our Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with independent type distribution $\mathbf{p}$ has two classes of players, $\mathcal{U}_1$ and $\mathcal{U}_2$:

- The class $\mathcal{U}_1$ consists of m players with type set $T_i = \{t_i, t'_i\}$, where $w(t_i) = 1, w(t'_i) = 0, p(i,t_i) = \frac{1}{\sqrt{m}}$ and $p(i,t'_i) = 1 - \frac{1}{\sqrt{m}}$ for all players $i \in \mathcal{U}_1$.
- The class $\mathcal{U}_2$ consists of $(\sqrt{m}-1)\cdot m$ players with type set $T_i = \{t_i\}$, where $w(t_i) = \frac{1}{\sqrt{m}}$ and $p(i,t_i) = 1$ all players $i \in \mathcal{U}_2$.

Consider the pure strategy profile σ' that assigns to each link one player from $\mathcal{U}_1$ and $\sqrt{m}-1$ players from $\mathcal{U}_2$. By analyzing the social cost of σ', we get

$$\begin{aligned}\mathsf{SC}_{\mathsf{MSP}}(\sigma',\Gamma) &\le 1 + (\sqrt{m}-1)\cdot\frac{1}{\sqrt{m}}\\ &< 2.\end{aligned}$$

Now consider the normal pure strategy profile σ where $\sqrt{m}$ players from $\mathcal{U}_1$ are assigned to each link $j \in [\sqrt{m}]$ and $\sqrt{m}$ players from $\mathcal{U}_2$ to each of the remaining $m - \sqrt{m}$ links. Clearly, σ is a normal pure Bayesian Nash equilibrium.

To show a lower bound on $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)$ we consider any link $j \in [\sqrt{m}]$. The actual load, say X_j, on link $j \in [\sqrt{m}]$ is a random variable which is a sum of $\sqrt{m}$ independent random variables with $\mathbb{E}(X_j) = 1$. Let $1 \le k \le \sqrt{m}, k \in \mathbb{N}$; the precise choice of k will be made later. Clearly,

$$\begin{aligned}\Pr(X_j \ge k) &\ge \Pr(X_j = k)\\ &= \binom{\sqrt{m}}{k}\cdot\left(\frac{1}{\sqrt{m}}\right)^k\cdot\left(1-\frac{1}{\sqrt{m}}\right)^{\sqrt{m}-k}\\ &\ge \binom{\sqrt{m}}{k}\cdot\left(\frac{1}{\sqrt{m}}\right)^k\cdot\left(1-\frac{1}{\sqrt{m}}\right)^{\sqrt{m}-1}\\ &\ge \binom{\sqrt{m}}{k}\cdot\left(\frac{1}{\sqrt{m}}\right)^k\cdot\frac{1}{\mathsf{e}} \qquad \text{(see [88, Lemma 2])}\\ &= \frac{\sqrt{m}\cdot\ldots\cdot(\sqrt{m}-k+1)}{\sqrt{m}^k}\cdot\frac{1}{k!}\cdot\frac{1}{\mathsf{e}}.\end{aligned}$$

Now, observe that $\frac{\sqrt{m}\cdot\ldots\cdot(\sqrt{m}-k+1)}{\sqrt{m}^k}$ is monotonically increasing in $\sqrt{m}$. Thus we get with $\sqrt{m} \geq k$,

$$\begin{aligned}\frac{\sqrt{m}\cdot\ldots\cdot(\sqrt{m}-k+1)}{\sqrt{m}^k} &\geq \frac{k\cdot\ldots\cdot(k-k+1)}{k^k}\\ &= \frac{k!}{k^k}.\end{aligned}$$

It follows that

$$\begin{aligned}\Pr(X_j \geq k) &\geq \frac{k!}{k^k}\cdot\frac{1}{k!}\cdot\frac{1}{\mathsf{e}}\\ &= \frac{1}{\mathsf{e}\cdot k^k},\end{aligned}$$

so that

$$\Pr(X_j < k) \leq 1 - \frac{1}{\mathsf{e}\cdot k^k}.$$

Now, since the actual loads $X_1,\ldots,X_{\sqrt{m}}$ are independent of each other, we have

$$\begin{aligned}\Pr((X_1 < k)\wedge\ldots\wedge(X_{\sqrt{m}} < k)) &= \prod_{j\in[\sqrt{m}]}\Pr(X_j < k)\\ &\leq \left(1-\frac{1}{\mathsf{e}\cdot k^k}\right)^{\sqrt{m}}\\ &\leq \mathsf{e}^{-\frac{1}{\mathsf{e}\cdot k^k}\cdot\sqrt{m}},\end{aligned}$$

where the last inequality holds since $1-x \leq \mathsf{e}^{-x}$ for all $x \in [0,1]$. Define now $\alpha > 0$ so that $\left(\frac{\alpha}{\mathsf{e}}\right)^\alpha = m$. Then, clearly[1], $\alpha = \Theta(\frac{\log m}{\log\log m})$. Choose $k = \frac{\alpha}{\mathsf{e}}$. We get that $k^k = (\frac{\alpha}{\mathsf{e}})^{\frac{\alpha}{\mathsf{e}}} = m^{\frac{1}{\mathsf{e}}}$ and thus

$$\begin{aligned}\Pr((X_1 < k)\wedge\ldots\wedge(X_{\sqrt{m}} < k)) &\leq \mathsf{e}^{-\frac{1}{\mathsf{e}\cdot k^k}\cdot\sqrt{m}}\\ &= \mathsf{e}^{-\frac{1}{\mathsf{e}}\cdot m^{\frac{1}{2}-\frac{1}{\mathsf{e}}}}\\ &\leq \frac{1}{m},\end{aligned}$$

[1] The Gamma factorial function Γ_F [54, Chapter 6] is for any integer $N \geq 1$ defined as $\Gamma_F(N+1) = N!$ and it is well known that $\Gamma_F^{-1}(N) = (\log N)/(\log\log N)\cdot(1+o(1))$. Thus Stiriling's formula $N! = (\frac{N}{\mathsf{e}})^N\cdot\sqrt{2\pi N}\cdot(1+o(1))$ can be used to get that $\left(\frac{\alpha}{\mathsf{e}}\right)^\alpha = m$ for $\alpha = \Theta(\log m/\log\log m)$.

where the last inequality holds since m is larger than 10^{15}. This implies that

$$\begin{aligned}
\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) &\geq \Pr((X_1 \geq k) \vee ... \vee (X_{\sqrt{m}} \geq k)) \cdot k \\
&= \big(1 - \Pr((X_1 < k) \wedge ... \wedge (X_{\sqrt{m}} < k))\big) \cdot k \\
&\geq \left(1 - \frac{1}{m}\right) \cdot \frac{\alpha}{\mathsf{e}} \\
&= \Theta\left(\frac{\log m}{\log \log m}\right).
\end{aligned}$$

Thus,

$$\begin{aligned}
\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} &\geq \frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{SC}_{\mathsf{MSP}}(\sigma',\Gamma)} \\
&= \Omega\left(\frac{\log m}{\log \log m}\right),
\end{aligned}$$

as needed. ■

8.4.2 Social Cost as Sum of Private Costs

In this section, we study the price of anarchy for social cost as the sum of private costs. In Theorem 8.7, we get that here fully mixed Bayesian Nash equilibria have worst social cost. This result is then used to prove an asymptotically tight bound on the price of anarchy (Theorem 8.8).

Proposition 8.4 states that the private cost of each player is maximized in a fully mixed Bayesian Nash equilibrium. Hence, we obtain:

Theorem 8.7: *Consider a Bayesian routing game* Γ *on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and a Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

$$\mathsf{SC}_{\mathsf{SUM}}(\mathbf{Q},\Gamma) \leq \mathsf{SC}_{\mathsf{SUM}}(\mathbf{F},\Gamma).$$

We now use Theorem 8.7 to prove an asymptotically tight bound on the price of anarchy for the case of identical links.

Theorem 8.8: *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links and an associated Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

$$\frac{\mathsf{SC}_{\mathsf{SUM}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{SUM}}(\Gamma)} \leq \frac{m+n-1}{m},$$

and this bound is tight up to a factor of $(1+\epsilon)$ *for any* $\epsilon > 0$*, even if* Γ *is a complete information routing game.*

Proof: By Theorem 8.7, it suffices to prove the upper bound for a fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Clearly,

$$\begin{aligned}\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F},\Gamma) &= \sum_{i\in[n]} \mathsf{PC}_i(\mathbf{F},\mathbf{p}) \\ &= \sum_{i\in[n]} \left(\frac{W}{m} + \frac{m-1}{m}\cdot W(i)\right) \quad \text{(by Proposition 8.3)} \\ &= \frac{nW}{m} + \frac{m-1}{m}\cdot W \\ &= \frac{m+n-1}{m}\cdot W.\end{aligned}$$

On the other hand, $\mathsf{PC}_i(\mathbf{Q},\mathbf{p}) \geq W(i)$ for any player $i \in [s]$ and any strategy profile $\mathbf{Q}$; hence,

$$\mathsf{OPT}_{\mathsf{SUM}}(\Gamma) \;\geq\; \sum_{i\in[n]} W(i) \;=\; W.$$

The upper bound follows.

We now prove that this upper bound is tight even for complete information routing games. To do so, we will prove that for any $\varepsilon > 0$, there is a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, T, 1)$ such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon)\cdot W$. We proceed by case analysis on the relation between n and m.

- Assume first that $n \leq m$. Let Γ_{CI} be an arbitrary complete information routing game with $n \leq m$. Then we can assign each player to a separate link which yields $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) = W$.

- Assume now that $n > m$. Define the complete information routing game Γ_{CI} as follows:

  > There are two sets of players $\mathcal{U}_1$, and $\mathcal{U}_2$. The set $\mathcal{U}_1$ consists of $n-m+1$ players with $w(t_i) = 1$ for all $i \in \mathcal{U}_1$, and $\mathcal{U}_2$ consists of $m-1$ players with $w(t_i) = k$ for all $i \in \mathcal{U}_2$ where $k \in \mathbb{N}$ is a constant to be determined later.

  For the (expected) total weight, we get

  $$W = n - m + 1 + (m-1)\cdot k.$$

  Let σ be the pure strategy profile that assigns all players from $\mathcal{U}_1$ to link m and each of the $m-1$ players from $\mathcal{U}_2$ separately to a link from $[m-1]$.

Thus,

$$
\begin{aligned}
&\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \\
&\quad \leq \mathsf{SC}_{\mathsf{SUM}}(\sigma, \Gamma_{\mathsf{CI}}) \\
&\quad = (n-m+1)^2 + (m-1)\cdot k \\
&\quad = \frac{(n-m+1)^2 + (m-1)\cdot k}{n-m+1+(m-1)\cdot k} \cdot W \\
&\quad = \frac{(n-m+1)\cdot(n-m) + (n-m+1) + (m-1)\cdot k}{n+(m-1)\cdot(k-1)} \cdot W \\
&\quad = \left(1 + \frac{(n-m)\cdot(n-m+1)}{n+(m-1)\cdot(k-1)}\right)\cdot W.
\end{aligned}
$$

Clearly, for any $\varepsilon > 0$, there is a $k \in \mathbb{N}$ such that $\frac{(n-m)\cdot(n-m+1)}{n+(m-1)\cdot(k-1)} \leq \varepsilon$. Hence, for any $\varepsilon > 0$, there is a complete information routing game Γ_{CI} such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon)\cdot W$. This completes the proof for the case $n > m$.

In all cases, there is a complete information routing game Γ_{CI} such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon)\cdot W$. Since $\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F}, \Gamma_{\mathsf{CI}}) = \frac{m+n-1}{m}\cdot W$, it follows that

$$
\frac{\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F}, \Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}})} \geq \frac{1}{1+\varepsilon}\cdot\frac{m+n-1}{m},
$$

as needed. ■

Berenbrink et al. [14] showed that the price of anarchy for complete information routing games (i.e., weighted congestion games on parallel links) and social cost as sum of private costs grows at least linearly with the number of players. In particular, they proved that $\frac{n}{5}$ is a lower bound on the price of anarchy. Theorem 8.8 implies that the price of anarchy increases at most linear with n and also shows the impact of the number of links.

Another interesting insight of Theorem 8.8 is that the price of anarchy does *not* increase if we allow incomplete information. This is not the case if social cost is defined as the maximum private cost, as we will see next.

8.4.3 Social Cost as Maximum Private Cost

In this section, we study the price of anarchy for social cost as the maximum private cost. As in the last section our asymptotic tight bounds on the price of anarchy (Theorem 8.10) are based on the fact that fully mixed Bayesian Nash equilibria have worst social cost (Theorem 8.9). We obtain the latter with Proposition 8.4:

Theorem 8.9: *Consider a Bayesian routing game* Γ *on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and a Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

$$\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma) \;\leq\; \mathsf{SC}_{\mathsf{MAX}}(\mathbf{F},\Gamma).$$

We now use Theorem 8.9 to prove asymptotically tight bounds on the price of anarchy for both Bayesian routing games and complete information routing games on identical links.

Theorem 8.10: *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links and an associated Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

(a) $\dfrac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \;\leq\; \dfrac{m+n-1}{m}$, *and*

(b) $\dfrac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \;\leq\; 2 - \dfrac{1}{m}$, *if* Γ *is a complete information routing game.*

The bound from (a) is tight up to a factor of $(1+\epsilon)$ *for any* $\epsilon > 0$ *and the bound from (b) is tight.*

Proof: Let $\mathbf{F}$ be a fully mixed Bayesian Nash equilibrium for Γ. Propositions 8.3 and 8.4 imply together that

$$\mathsf{PC}_i(\mathbf{Q},\mathbf{p}) \;\leq\; \mathsf{PC}_i(\mathbf{F},\mathbf{p}) \;=\; \frac{W}{m} + \frac{m-1}{m}\cdot W(i), \tag{8.4}$$

for each player $i \in [n]$. We now prove the two parts (a) and (b) of the theorem.

Part (a), upper bound: Clearly, for any strategy profile $\mathbf{Q}'$ and for any player $i \in [n]$, $\mathsf{PC}_i(\mathbf{Q}',\mathbf{p}) \geq W(i)$; hence, $\sum_{i\in[n]} \mathsf{PC}_i(\mathbf{Q}',\mathbf{p}) \geq W$. This implies that

$$\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \;\geq\; \frac{W}{n}. \tag{8.5}$$

Clearly, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \geq W(i)$ for all $i \in [n]$. Fix any player $i \in [n]$. By (8.4) and (8.5),

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{Q},\mathbf{p}) &\leq \frac{W}{m} + \frac{m-1}{m}\cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma)\\
&\leq \frac{n}{m}\cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma) + \frac{m-1}{m}\cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma)\\
&= \frac{m+n-1}{m}\cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma).
\end{aligned}$$

and the upper bound follows.

Part (a), lower bound: Fix any arbitrary $k, a, r \in \mathbb{N}$, which will be determined later. Consider the Bayesian routing game $\Gamma_{k,a,r} = (n, m, \mathbf{1}, T, \mathbf{p})$ with independent type distribution and $n = k \cdot (m-1)$ players. Each player $i \in [n]$ has type set $T_i = \{t_i, t'_i\}$ with weights $w(t_i) = 1$, $w(t'_i) = a \cdot r$ and probabilities $p(i, t_i) = 1 - \frac{1}{a}$, $p(i, t'_i) = \frac{1}{a}$. Clearly, for player $i \in [n]$, $W(i) = r + 1 - \frac{1}{a}$.

Define a pure strategy profile σ that assigns all types t'_i, $i \in [n]$, of weight 1 to link m. The types t_i, $i \in [n]$, are evenly distributed among the links in $[m-1]$; so, σ assigns exactly k of these types to each link in $[m-1]$. Now for each player $i \in [n]$,

$$\begin{aligned}
&\mathsf{PC}_i(\sigma, \mathbf{p}) \\
&= \left(1 - \frac{1}{a}\right) \cdot \left(1 + (k-1) \cdot \left(1 - \frac{1}{a}\right)\right) + \frac{1}{a} \cdot ((n-1) \cdot r + r \cdot a) \\
&= \left(1 - \frac{1}{a}\right) \cdot \left(\frac{1}{a} + k \cdot \left(1 - \frac{1}{a}\right)\right) + r \cdot \left(\frac{(n-1)}{a} + 1\right);
\end{aligned}$$

so, for any $\epsilon' > 0$, there is a sufficiently large a such that for each player $i \in [n]$,

$$\mathsf{PC}_i(\sigma, \mathbf{p}) \leq (k + r) \cdot (1 + \epsilon').$$

Hence, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r}) \leq (k+r)(1+\epsilon')$. Fix now any fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Proposition 8.3 implies that for each player $i \in [n]$,

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{F}, \mathbf{p}) &= \left(1 + \frac{n-1}{m}\right) \cdot W(i) \\
&= \frac{m+n-1}{m} \cdot \left(r + 1 - \frac{1}{a}\right).
\end{aligned}$$

Thus $\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r}) = \frac{m+n-1}{m} \cdot \left(r + 1 - \frac{1}{a}\right)$ and we can conclude that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r})}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r})} \geq \frac{(r + 1 - \frac{1}{a})}{(k+r)(1+\epsilon')} \cdot \frac{m+n-1}{m};$$

so, for any $\epsilon > \epsilon'$, there is a sufficiently large r such that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r})}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r})} \geq \frac{m+n-1}{m} \cdot \frac{1}{1+\epsilon}.$$

This proves that the upper bound shown before is tight up to a factor of $(1+\epsilon)$.

Part (b), upper bound: Consider for the upper bound a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t_1, \ldots, t_n)\}, 1)$. Here, $W(i) = w(t_i)$

for all $i \in [n]$. Clearly, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) \geq W(i)$ for all $i \in [n]$ and $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) \geq \frac{W}{m}$. By Equation (8.4),

$$\begin{aligned}\mathsf{PC}_i(\mathbf{Q}, \mathbf{p}) &\leq \frac{W}{m} + \frac{m-1}{m} \cdot W(i)\\ &\leq \mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) + \frac{m-1}{m} \cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}})\\ &= \left(2 - \frac{1}{m}\right) \cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}),\end{aligned}$$

so that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q}, \Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \leq 2 - \frac{1}{m}$$

and hence the upper bound follows.

Part (b), lower bound: Consider for the lower bound the complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t_1, \ldots, t_n)\}, 1)$ with $n = m$ and $w(t_1) = \ldots = w(t_n) = 1$. Clearly, $W(i) = w(t_i) = 1$ for all $i \in [n]$, $W = m$ and $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) = 1$. Now, for the fully mixed Nash equilibrium $\mathbf{F}$ and any player $i \in [n]$, by Equation (8.4),

$$\begin{aligned}\mathsf{PC}_i(\mathbf{F}, \mathbf{p}) &= \frac{W}{m} + \frac{m-1}{m} \cdot W(i)\\ &= (2 - \frac{1}{m}) \cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}),\end{aligned}$$

so that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} = 2 - \frac{1}{m}. \qquad \blacksquare$$

8.5 Conclusion and Directions for Further Research

In this chapter, we studied Bayesian routing games with incomplete information. With the help of a potential function we were able to prove that every Bayesian routing game possesses a pure Bayesian Nash equilibrium. For games with identical links and independent type distribution, we described a polynomial time algorithm that is able to compute such a pure Bayesian Nash equilibrium. We showed that our algorithm does not work for Bayesian routing games with correlated type distribution or related links. For these cases it is an open problem whether a pure Bayesian Nash equilibrium can be computed in polynomial time.

Our study of the structural properties of fully mixed Bayesian Nash equilibria in the model of identical links revealed that the private costs of all players are maximized in a fully mixed Bayesian Nash equilibrium. This interesting insight enabled us to prove (asymptotic) tight bounds on the price of anarchy for identical link games if the social cost measure is given by the average or the maximum private cost of a player. Furthermore, we considered the price of anarchy for identical link games if instead social cost is described by the expected maximum latency on a link. Here, we established lower and upper bounds for different special cases but the exact price of anarchy has been left open. In this chapter, we focused our study on the price of anarchy only to identical link games. The next logical step is to extend the given results to related links.

CHAPTER 9

Bottleneck Games with Splittable Traffic

9.1 Introduction

While Wardrop games are appropriate to model road traffic systems, the basic assumption of drivers minimizing their travel time (more generally called path latency) is not reasonable in all networks. For instance, in communication networks such as the Internet, providers of streaming content would try to maximize the *throughput* to their clients whereas the transmission time is of lesser concern. More generally, for a routing path in such a network, one would be interested in the *maximum* latency of all its edges—in other words, the *latency of the bottleneck*—as it is inversely proportional to the achievable throughput on that very path.

From a purely mathematical perspective, the bottleneck latency of a path corresponds to the ∞-norm of the vector of edge latencies whereas the sum of edge latencies—that was of interest in Wardrop games—equals the 1-norm. For a broader discussion of when the ∞-norm should be used, confer also Banner and Orda [11]. They mention, for instance, that the ∞-norm is appropriate to model wireless networks where each node has a limited transmission energy.

As another motivation, a celebrated result by Leighton et al. [61] implies that the bottleneck latency is also of interest in settings with individual traffic: Their result states that individual packets can be routed in time $\mathcal{O}$(congestion + dilation) when the paths for the packets are given in advance. Here dilation denotes the maximum length of a path and congestion denotes the maximum number of paths sharing a common edge.

We address the bottleneck latency scenarios described in the last paragraphs by studying the *bottleneck games with splittable traffic* that were formally introduced in Section 3.4. Similar to Wardrop games, one could for such a game think of infinitely many selfish players each controlling a negligible amount of traffic. However, their objective is now to choose a path such that their experienced bottleneck latency is at minimum. Likewise, we define a Wardrop equilibrium in a game of our new model as a traffic distribution where no fraction of the traffic assigned to some path, however small, can decrease its bottleneck latency by unilaterally switching to another path.

A part of our results focuses on bottleneck games with splittable traffic and *M/M/1 latency functions*. These latency functions arise in queuing theory as the expected latency of queues with a Poisson arrival process and an exponentially distributed service time [57, 79]. They are used in networking theory to model packet-switched networks. Here, a packet that starts at its entry node in the network or arrives at an intermediate node on its way to the destination is stored in a queue. It can leave the queue as soon as the next link on the path of the packet becomes available [16, 87].

9.1.1 Contribution

Our investigations are two-fold: First, we study general properties of bottleneck games with splittable traffic such as existence and uniqueness of Wardrop equilibria and dependence of both the price of anarchy and stability of the network topology. Most of our results here are based on properties of maximum flows and minimum cuts[1]. In the second part we prove an exact expression for the price of stability for bottleneck games with splittable traffic and M/M/1 latency functions on parallel links. In detail, our main findings are:

- General results for bottleneck games with splittable traffic:
  - We show that a bottleneck game with splittable traffic possesses a Wardrop equilibrium with finite social cost if the traffic is smaller than the network capacity.
  - For bottleneck games with splittable traffic on series parallel graphs we prove the social cost of Wardrop equilibria to be unique. By contrast, we also show that for any graph whose subgraph induced by all simple origin-destination paths is not series parallel, there exists a game having equilibria with different social cost.
  - We show that the price of stability for bottleneck games with splittable traffic is independent of the network topology, i.e., the worst-case ratio, over all instances, between the best Nash equilibrium and an optimum is attained on parallel links. (See Section 9.2.3 for a comparison with a similar result by Cole et al. [22].)

[1] See Section 2.5 for a brief description of maximum flows and minimum cuts.

- Bottleneck games with splittable traffic and M/M/1 latency functions:
We prove that the expression

$$\frac{m \cdot \frac{r}{c_{\min}}}{\frac{r}{c_{\min}} + 2 \cdot (m-1) \cdot \left(\sqrt{\frac{r}{c_{\min}} + 1} - 1\right)} \tag{9.1}$$

describes the exact price of stability for games with M/M/1 latency functions, minimum edge capacity $c_{\min}$, and traffic r on m parallel links.

Since bottleneck games with splittable traffic on parallel links are Wardrop games, it is possible to draw interesting conclusions based on our result that the price of stability for bottleneck games with splittable traffic is independent of the network topology. To be more precise, it follows that the price of stability for bottleneck games with splittable traffic on arbitrary graphs corresponds to the price of anarchy for Wardrop games on parallel links. This can be used when latency functions are restricted to polynomials where results of Roughgarden [82] can be applied, and also for the class of M/M/1 latency functions where our expression (9.1) describes the price of stability for bottleneck games with splittable traffic on arbitrary graphs.

9.1.2 Related Work

The related work that is relevant for this chapter studies Wardrop games (Section 4.4), bottleneck games (Section 4.5), and finite splittable routing games (Section 4.6).

9.1.3 Road Map

The rest of this chapter is organized as follows. We study bottleneck games with splittable traffic and general latency functions in Section 9.2, whereas we restrict ourselves to M/M/1 latency functions in Section 9.3. We conclude in Section 9.4.

9.2 Games with General Latency Functions

For bottleneck games with splittable traffic we will prove the existence of Wardrop equilibria (Section 9.2.1), study the (non-)uniqueness of equilibria social cost (Section 9.2.2), and show that the price of stability is independent of the network topology (Section 9.2.3). Before we focus on these results we want to make a simple observation which shows that not restricting the paths of $\mathcal{P}_{v_o v_d}$ to be simple would yield a different game (which is not the case for the classic Wardrop games where cycle paths could be omitted without loss of generality).

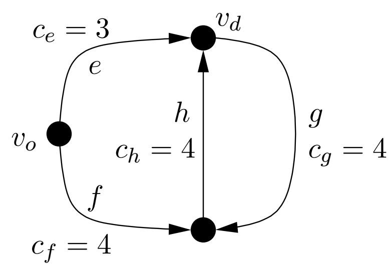

Figure 9.1: Graph used in the proof of Observation 9.1

Observation 9.1: *There is a bottleneck game Γ with splittable traffic that has a unique Wardrop equilibrium if only simple paths are allowed and two different Wardrop equilibria if also non-simple paths are allowed.*

Proof: Consider the bottleneck game $\Gamma = (G, v_o, v_d, (c_e)_{e \in E}, r)$ with splittable traffic and M/M/1 latency functions whose graph is shown in Figure 9.1. Edge e has capacity 3 whereas all others are of capacity 4. The traffic to be routed is $r = 3$. Since only simple paths are allowed, there are only two possible paths from v_o to v_d. Let $\mathbf{x}$ denote the strategy profile where the amount of traffic using path (e) is 1 and the amount on (f, h) is 2. It is easy to see that $\mathbf{x}$ is the unique Wardrop equilibrium. If paths were allowed that are not simple then the strategy profile where the amount of traffic on (e, g, h) is 2 and the amount on (f, h) is 1 would also be a Wardrop equilibrium. ■

9.2.1 Existence of Wardrop Equilibria

Existence of Wardrop equilibria in bottleneck games with splittable traffic can be established by employing the general result of [86] (for a proof using more elementary maths, see [78]). To illustrate the connection to maximum flows, however, we start by giving a proof that makes use of the max-flow min-cut theorem (see, e.g., [3]). The construction described in the proof will also be used in the proof of Theorem 9.5. Obviously, the only interesting case is the traffic to be routed being smaller than the capacity of the network.

Theorem 9.1: *Let $\Gamma = (G, v_o, v_d, (f_e)_{e \in E}, r)$ be a bottleneck game with splittable traffic where $r < C(G, v_o, v_d, (f_e)_{e \in E})$. Then Γ possesses a Wardrop equilibrium of finite social cost.*

Proof: We give a short outline of the proof first: Restricting the latency of any used path in $\mathcal{P}_{v_o v_d}$ to $y \geq 0$ induces maximum flow instances with edge capacities each within an interval $[k_e^{\min}(y), k_e^{\max}(y)] \subseteq \mathbb{R}_0^+$. Any solution to one of these instances is a Wardrop equilibrium of a game $(G, v_o, v_d, (f_e)_{e \in E}, r')$ where $r' \in R(y)$ and $R(y) \subseteq \mathbb{R}_0^+$ is an interval. The theorem follows as for any $r' \in R(y)$ there exists a matching maximum flow instance (and hence a Wardrop equilibrium) and furthermore $\bigcup_{y \in \mathbb{R}_0^+} R(y) = [0, C(G, v_o, v_d, (f_e)_{e\in E}))$.

We now present the proof in detail. Let $\Gamma = (G, v_o, v_d, (f_e)_{e\in E}, r)$ be as stated in the theorem. W.l.o.g., we can assume for any $e \in E$ that $f_e(u) \to \infty$ for $u \to \infty$. (Since the traffic $r \in \mathbb{R}_0^+$ is finite and therefore the definition of f_e on (r, ∞) has no effect.) For every edge $e \in E$ we will now introduce the functions $k_e^{\min}$ and $k_e^{\max}$ where for a $y \geq f_e(0)$ the values $k_e^{\min}(y)$ and $k_e^{\max}(y)$ will be the minimum and maximum load that we can have on edge e to get a latency of y. For an edge $e \in E$ and $\text{ext} \in \{\min, \max\}$, let $k_e^{\text{ext}} : \mathbb{R}_0^+ \to \mathbb{R}_0^+$,

$$k_e^{\text{ext}}(y) \quad = \quad \begin{cases} 0 & \text{if } y < f_e(0) \\ \text{ext}\, f^{-1}(\{y\}) & \text{if } y \geq f_e(0). \end{cases}$$

Note that $k_e^{\min}$ and $k_e^{\max}$ are well-defined functions because non-empty closed subsets of $\mathbb{R}$ that are bounded below (resp., above) have a minimum (resp., maximum). Observe that for any edge $e \in E$, a latency $y \in \mathbb{R}_0^+$, and an edge flow $u \in [k_e^{\min}(y), k_e^{\max}(y)]$, it holds that $f_e(u) = y$ if $y \geq f_e(0)$ and $f_e(u) = f_e(0) > y$ otherwise. Obviously, the functions $k_e^{\min}$ and $k_e^{\max}$ are non-decreasing. From elementary analysis, we have that $k_e^{\min}$ is lower semi-continuous whereas $k_e^{\max}$ is upper semi-continuous.

Here we will only prove that $k_e^{\max}$ is upper semi-continuous. (The lower semi-continuity for $k_e^{\min}$ follows with similar arguments.) Assume, by the way of contradiction, that there are $y \geq f_e(0)$ and $\delta > 0$ such that for every $\varepsilon > 0$ it holds that $k^{\max}(y + \varepsilon) \geq k^{\max}(y) + \delta$. By definition of $k^{\max}$, we have that $f(k^{\max}(y) + \delta) = y + \varepsilon_y$ for an $\varepsilon_y > 0$. Since f is non-decreasing, we have for all $\varepsilon > 0$ that

$$y + \varepsilon \;=\; f(k^{\max}(y + \varepsilon)) \;\geq\; f(k^{\max}(y) + \delta) \;=\; y + \varepsilon_y,$$

i.e., $\varepsilon \geq \varepsilon_y$ which is a contradiction since $\varepsilon_y > 0$.

Define now $\ell((k_e)_{e\in E})$, $\ell : (\mathbb{R}_0^+)^{|E|} \to \mathbb{R}_0^+$, as the maximum flow for the maximum flow problem instance $(G, v_o, v_d, (k_e)_{e\in E})$, i.e., $\ell((k_e)_{e\in E})$ is the maximum amount of flow possible if k_e is the capacity restriction for e. ℓ is continuous because if $k = (k_e)_{e\in E}$ and $k' = (k'_e)_{e\in E}$ are capacity vectors with $k_e - \epsilon \leq k'_e \leq k_e + \epsilon$ for all $e \in E$ then

$$\ell(k) - \epsilon \cdot |E| \;\leq\; \ell(k') \;\leq\; \ell(k) + \epsilon \cdot |E|.$$

Finally, for $\text{ext} \in \{\max, \min\}$ set $r^{\text{ext}}(y) = \ell((k_e^{\text{ext}}(y))_{e\in E})$, $r^{\text{ext}} : \mathbb{R}_0^+ \to \mathbb{R}_0^+$. Hence, $r^{\min}(y)$ is the maximum flow we achieve if we select for each edge e as capacity $k_e^{\min}(y)$ either the smallest load that causes latency y or 0 if even load 0 leads to a latency of at least y. In almost the same manner $r^{\max}(y)$ is the maximum flow we achieve if we select for each edge e as capacity $k_e^{\max}(y)$ either the largest load that causes latency y or 0 if even load 0 leads to a latency greater than y.

Now fix $y \in \mathbb{R}_0^+$ and $(k_e)_{e\in E}$ where $k_e \in [k_e^{\min}(y), k_e^{\max}(y)]$ for all $e \in E$. Consider the maximum flow problem on the instance $(G, v_o, v_d, (k_e)_{e\in E})$ for which $H = (h_e)_{e\in E}$ is a maximum flow. W.l.o.g. there is no cycle in the graph G whose edges e all have a non-zero flow $h_e > 0$. (By the flow decomposition theorem [3] it is possible to remove such cycle-flows. The outcome is again a maximum flow.) We now perform a flow decomposition to get a strategy profile $\mathbf{x}$ for the bottleneck game with splittable traffic $(G, v_o, v_d, (f_e)_{e\in E}, \ell((k_e)_{e\in E}))$ where $\delta_e(\mathbf{x}) = h_e$ for all edges $e \in E$.

We will now show that $\mathbf{x}$ is a Wardrop equilibrium: By the max-flow min-cut theorem (see Section 2.5) we have that there exists a cut $D(S,T) \subseteq E$ with $h_e = k_e$ for all $e \in D(S,T)$. Consider an arbitrary path $P \in \mathcal{P}_{v_o v_d}$: Assume first that $x_P > 0$, hence we have for all $e \in P$ that $k_e > 0$ and $y \geq f_e(0)$. Obviously, P contains an edge $\hat{e} \in D(S,T)$, i.e., an edge $\hat{e}$ with

$$f_{\hat{e}}(\delta_{\hat{e}}(\mathbf{x})) \;=\; f_{\hat{e}}(h_{\hat{e}}) \;=\; f_{\hat{e}}(k_{\hat{e}}) \;=\; y.$$

Furthermore, $h_e \leq k_e$ for all edges $e \in P$ and therefore

$$f_e(\delta_e(\mathbf{x})) \;=\; f_e(h_e) \;\leq\; f_e(k_e) \;=\; y.$$

If, on the other hand, $x_P = 0$, then P has to contain an edge $e \in D(S,T)$ with $h_e = k_e$ and

$$f_e(\delta_e(\mathbf{x})) \;=\; f_e(h_e) \;=\; f_e(k_e) \;\geq\; y.$$

We eventually get that each path with non-zero flow has latency y, whereas each path without flow has latency at least y. Hence, $\mathbf{x}$ is an equilibrium.

Finally, since $r^{\min}(0) = 0$, $r^{\max}(y) \to C(G, v_o, v_d, (f_e)_{e\in E})$ for $y \to \infty$, $r^{\min}$ is lower semi-continuous, and $r^{\max}$ is upper semi-continuous, there has to be a $y \in \mathbb{R}_0^+$ with $r^{\min}(y) \leq r$ and $r^{\max}(y) \geq r$. Continuity of ℓ and the preceding paragraph then assure the existence of a Wardrop equilibrium for Γ. ■

9.2.2 (Non-)Uniqueness Results about Social Cost of Equilibria

We will show in this section that different equilibria for a bottleneck game with splittable traffic on a series parallel graph have the same social cost. The proof for this result employs a technique based on what we define as *strong cuts*.

Definition 9.1: *Let Γ be a bottleneck game with splittable traffic on a series parallel graph $G = (V,E)$ and let $\mathbf{x}$ be a Wardrop equilibrium for Γ. Then $D \subseteq E$ is called* strong cut *with respect to Γ and $\mathbf{x}$ if*

1. *each path $P \in \mathcal{P}_{v_o v_d}$ contains exactly one edge that belongs to D, and*

2. *$f_e(\delta_e(\mathbf{x})) \geq l(\mathbf{x})$ for all edges $e \in D$.*

Observe that, given a strong cut D with respect to Γ and an equilibrium $\mathbf{x}$, all edges $e \in D$ with $\delta_e(\mathbf{x}) > 0$ have latency $l(\mathbf{x})$ whereas all other edges $e \in D$ with $\delta_e(\mathbf{x}) = 0$ have latency at least $l(\mathbf{x})$. Before making use of the crucial properties of strong cuts, we need to ensure their existence.

Theorem 9.2: *Let Γ be a bottleneck game with splittable traffic on a series parallel graph and let $\mathbf{x}$ be a Wardrop equilibrium for Γ. Then a strong cut with respect to Γ and $\mathbf{x}$ exists.*

Proof: The proof is by structural induction over all series parallel graphs. Our *induction hypothesis* is that every series parallel graph G with terminals (v_o, v_d) has the following property: For any bottleneck game Γ with splittable traffic on G and all Wardrop equilibria of Γ, there is a strong cut.

The only *base case* to verify consists of the graph with two nodes v_o, v_d solely connected by the edge e. Obviously, for any game Γ on this graph, $\{e\}$ is a strong cut with respect to Γ and its trivial equilibrium. For the *induction step*, consider any arbitrary series parallel graph $G = (V, E)$ with terminals (v_o, v_d). Furthermore, assume that G is a parallel or series connection of two series parallel graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ with terminals (v_o^1, v_d^1) and (v_o^2, v_d^2), respectively, and both G_1 and G_2 fulfill the induction hypothesis. To prove the induction step, we then have to show that G fulfills the induction hypothesis, too. Thus, let $\Gamma = (G, v_o, v_d, (f_e)_{e \in E}, r)$ be an arbitrary game on G and $\mathbf{x}$ be an arbitrary Wardrop equilibrium for Γ and consider the two cases:

Parallel Connection: In this case $v_o = v_o^1 = v_o^2$ and $v_d = v_d^1 = v_d^2$. Set

$$r_1 = \sum_{P \in \mathcal{P}_{v_o^1 v_d^1}} x_P, \qquad r_2 = \sum_{P \in \mathcal{P}_{v_o^2 v_d^2}} x_P$$

where $\mathcal{P}_{v_o^1 v_d^1}$ and $\mathcal{P}_{v_o^2 v_d^2}$ are meant to only contain paths from G_1 and G_2, respectively. Obviously, $r_1 + r_2 = r$ and the two games $\Gamma_1 = (G_1, v_o^1, v_d^1, (f_e)_{e \in E_1}, r_1)$ and $\Gamma_2 = (G_2, v_o^2, v_d^2, (f_e)_{e \in E_2}, r_2)$ have Wardrop equilibria $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ where $x_P^{(1)} = x_P$ for all paths P with edges in E_1 and $x_P^{(2)} = x_P$ for all paths P with edges in E_2. It follows by the induction hypothesis that there are strong cuts D_1 and D_2 with respect to $G_1, \mathbf{x}^{(1)}$ and $G_2, \mathbf{x}^{(2)}$ that we can use to get a strong cut $D = D_1 \cup D_2$ for Γ and its equilibrium $\mathbf{x}$.

Series Connection: In this case $v_o = v_o^1$, $v_d^1 = v_o^2$, $v_d^2 = v_d$. Consider the games $\Gamma_1 = (G_1, v_o^1, v_d^1, (f_e)_{e \in E_1}, r)$ and $\Gamma_2 = (G_2, v_o^2, v_d^2, (f_e)_{e \in E_2}, r)$. Obviously, $\mathbf{x}$ induces strategy profiles $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ for Γ_1 and Γ_2, respectively. At least one of $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ is a Wardrop equilibrium. Otherwise, there would be a path $P \in \mathcal{P}_{v_o v_d}$ with non-zero flow on which the latency is larger than on another path $R \in \mathcal{P}_{v_o v_d}$, and $\mathbf{x}$ cannot be a Wardrop equilibrium. If $\mathbf{x}^{(1)}$ is an equilibrium we set $D = D_1$ and $D = D_2$ otherwise. In either case, D is a strong cut for Γ and its equilibrium $\mathbf{x}$. ∎

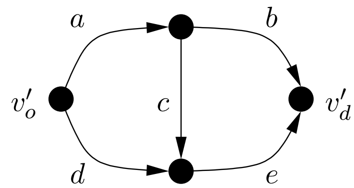

Figure 9.2: Braess paradox graph

We now use strong cuts in the proof of the next theorem to show that all Wardrop equilibria of a bottleneck game with splittable traffic on a series parallel graph have the same social cost. Obviously, this implies that the price of stability does not differ from the price of anarchy for this class of games.

Theorem 9.3: *Let Γ be a bottleneck game with splittable traffic on a series parallel graph, let $\hat{\mathbf{x}}$ and $\mathbf{x}$ be two Wardrop equilibria for Γ. Then $\mathsf{SC}(\hat{\mathbf{x}}, \Gamma) = \mathsf{SC}(\mathbf{x}, \Gamma)$.*

Proof: The proof is by contradiction. Assume, by the way of contradiction, that two different Wardrop equilibria $\hat{\mathbf{x}}$ and $\mathbf{x}$ for $\Gamma = (G, v_o, v_d, (f_e)_{e\in E}, r)$ are given such that $\mathsf{SC}(\hat{\mathbf{x}}, \Gamma) < \mathsf{SC}(\mathbf{x}, \Gamma)$. Clearly, $l(\hat{\mathbf{x}}) < l(\mathbf{x})$. Let D be a strong cut with respect to Γ and $\mathbf{x}$. Consider an edge $e \in D$ with $\delta_e(\hat{\mathbf{x}}) > 0$. Since $\hat{\mathbf{x}}$ is a Wardrop equilibrium and e is an edge of the strong cut D with respect to Γ and $\mathbf{x}$, we get that

$$f_e(\delta_e(\hat{\mathbf{x}})) \ \leq \ l(\hat{\mathbf{x}}) \ < \ l(\mathbf{x}) \ \leq \ f_e(\delta_e(\mathbf{x})),$$

which implies $\delta_e(\hat{\mathbf{x}}) < \delta_e(\mathbf{x})$ because f_e is non-decreasing. If instead an edge $e \in D$ with $\delta_e(\hat{\mathbf{x}}) = 0$ is considered we trivially obtain that $\delta_e(\hat{\mathbf{x}}) \leq \delta_e(\mathbf{x})$. Together we get that

$$r \ = \ \sum_{e\in D} \delta_e(\hat{\mathbf{x}}) \ < \ \sum_{e\in D} \delta_e(\mathbf{x}) \ = \ r,$$

which is a contradiction. ■

We will now consider graphs that are not series parallel. We start by observing that there is a game on such a graph that has equilibria of different social cost.

Lemma 9.1: *There exists a bottleneck game Γ' with splittable traffic possessing Wardrop equilibria of different social cost.*

Proof: Consider the bottleneck game $\Gamma' = (G', v'_o, v'_d, (c_e)_{e\in E'}, r')$ with splittable traffic and M/M/1 latency functions where $G' = (V', E')$ is the Braess paradox graph as shown in Figure 9.2. All edges $e \in E'$ have capacity $c_e = r'$, i.e., the capacities equal the amount of traffic to be routed.

The strategy profile $\mathbf{x} = (0, r', 0)$ that only uses the "zigzag" path (a, c, e) is a Wardrop equilibrium of social cost $\mathsf{SC}(\mathbf{x}, \Gamma') = \infty$. However, the profile

$\hat{\mathbf{x}} = (\frac{r'}{2}, 0, \frac{r'}{2})$ that splits the traffic evenly on the "upper" and "lower" paths (a, b) and (d, e) is also an equilibrium with $\mathsf{SC}(\hat{\mathbf{x}}, \Gamma') = 2$. Hence, $\mathbf{x}$ and $\hat{\mathbf{x}}$ are of different social cost. ■

We will now use Lemma 9.1 to get a stronger result that does not restrict to the Braess paradox graph. Instead this stronger result considers all graphs G that are of interest. Recall that whenever traffic is sent through a graph G only edges that are on a simple path from v_o to v_d can be used. So the same equilibria are obtained when playing the game not on G but on the maximum subgraph of G containing only edges that are on a simple path from v_o to v_d. This idea is captured by the following definition.

Definition 9.2: *A directed multigraph $G = (V, E)$ without isolated vertices where $v_o, v_d \in V$, $v_o \neq v_d$, is called* strongly (v_o, v_d)-connected *if every edge $e \in E$ is contained in a simple path from v_o to v_d.*

In [68] we use a result [27, 89] stating that an acyclic strongly (v_o, v_d)-connected graph G is series parallel if and only if the Braess paradox graph is not a minor of G. Based on this result it is possible to simulate the game Γ' given in Lemma 9.1 on any strongly (v_o, v_d)-connected graph that is not series parallel. Thus we get that:

Theorem 9.4 (Mazalov, Monien, Schoppmann, Tiemann [68]): *Let G be a strongly (v_o, v_d)-connected graph that is not series parallel. Then there exists a bottleneck game $\Gamma = (G, v_o, v_d, (f_e)_{e \in E}, r)$ with splittable traffic possessing Wardrop equilibria of different social cost.*

9.2.3 Price of Stability

In this section, we will show that the price of stability for bottleneck games with splittable traffic and latency functions from an arbitrary set of latency functions $\mathcal{F}$ is the same on general graphs as on parallel links, i.e., $\mathsf{PoS}(\mathscr{G}(\mathcal{F})) = \mathsf{PoS}(\mathscr{P}(\mathcal{F}))$. To do so, we will show that given a game Γ on a general graph with latency functions from $\mathcal{F}$ there exists a game Γ' on parallel links with latency functions from $\mathcal{F}$ and Wardrop equilibria $\mathbf{x}$ for Γ and $\hat{\mathbf{x}}$ for Γ' such that $\frac{\mathsf{SC}(\mathbf{x},\Gamma)}{\mathsf{OPT}(\Gamma)} \leq \frac{\mathsf{SC}(\hat{\mathbf{x}},\Gamma')}{\mathsf{OPT}(\Gamma')}$.

We assume that Cole et al. [22] proved a very similar result to establish their Theorem 4.6 (whose proof they had to omit due to lack of space). Since we need a rather technical formulation for our result on the price of stability for games with M/M/1 latency functions, we give the following Theorem 9.5.

Theorem 9.5: *Let $\Gamma = (G, v_o, v_d, (f_e)_{e \in E}, r)$, $G = (V, E)$, be a bottleneck game with splittable traffic where $r < C(G, v_o, v_d, (f_e)_{e \in E})$. Then there exist*

- *a bottleneck game $\Gamma' = (G', v'_o, v'_d, (f'_{e'})_{e' \in E'}, r)$ with splittable traffic on parallel links $G' = (V', E')$ where $|E'| \leq |E|$ and for each $e' \in E'$ there is an edge $e \in E$ such that $f'_{e'} = f_e$ and*

- *Wardrop equilibria* $\mathbf{x}$ *for* Γ *and* $\hat{\mathbf{x}}$ *for* Γ',

such that $\frac{\mathsf{SC}(\mathbf{x},\Gamma)}{\mathsf{OPT}(\Gamma)} \leq \frac{\mathsf{SC}(\hat{\mathbf{x}},\Gamma')}{\mathsf{OPT}(\Gamma')}$.

Proof: Let $\mathbf{x}$ be the equilibrium for Γ whose existence we proved in Theorem 9.1. Recall that $\mathbf{x}$ corresponds to a maximum flow $H = (h_e)_{e \in E}$ of a maximum flow problem instance $(G, v_o, v_d, (k_e)_{e \in E})$. Denote again by $D(S,T) \subseteq E$ the minimum cut and assume, w.l.o.g., that H does not have cycle-flows.

We now construct a bottleneck game Γ' with splittable traffic on $|D(S,T)|$ parallel links where the edges of Γ' have the same latency functions as the edges in $D(S,T)$. Set $\Gamma' = (G', v_o', v_d', (f_e')_{e \in E'}, r)$ where $G' = (V', E')$ and $V' = \{v_o', v_d'\}$. The set E' consists of $|D(S,T)|$ edges from v_o' to v_d' such that there is a bijection $g : D(S,T) \to E'$ with $f'_{g(e)} = f_e$ for all $e \in D(S,T)$.

The load values $(\delta_e(\mathbf{x}))_{e \in D(S,T)}$ can be used to define a strategy profile $\hat{\mathbf{x}} = (\hat{x}_{e'})_{e' \in E'}$ for Γ', where $\hat{x}_{g(e)} = \delta_e(\mathbf{x})$ for all $e \in D(S,T)$. We will show that $\hat{\mathbf{x}}$ is a Wardrop equilibrium for Γ'. Remember that all edges $e \in D(S,T)$ of the minimum cut are saturated, i.e., $h_e = k_e$. If, on the one hand, an edge $e \in D(S,T)$ with $0 = \delta_e(\mathbf{x}) = h_e = k_e$ is considered we have that $f_e(0) \geq y$ (see proof of Theorem 9.1). Thus all edges $g(e)$, $e \in D(S,T)$, without load, i.e., $\hat{x}_{g(e)} = \delta_e(\mathbf{x}) = 0$, are of latency

$$f'_{g(e)}(\hat{x}_{g(e)}) = f'_{g(e)}(0) = f_e(0) \geq y.$$

If, on the other hand, an edge $e \in D(S,T)$ with $0 < \delta_e(\mathbf{x}) = h_e = k_e$ is considered we have that $f_e(\delta_e(\mathbf{x})) = f_e(h_e) = f_e(k_e) = y$ (see proof of Theorem 9.1). Therefore, each edge $g(e)$, $e \in D(S,T)$, with non-zero load, i.e., $\hat{x}_{g(e)} = \delta_e(\mathbf{x}) \neq 0$, is of latency

$$f'_{g(e)}(\hat{x}_{g(e)}) = f_e(\delta_e(\mathbf{x})) = y,$$

and $\hat{\mathbf{x}}$ is a Wardrop equilibrium with $l(\hat{\mathbf{x}}) = y = l(\mathbf{x})$. Obviously, $\mathsf{SC}(\mathbf{x},\Gamma) = \mathsf{SC}(\hat{\mathbf{x}},\Gamma')$.

Observe that in Γ each path $P \in \mathcal{P}_{v_o v_d}$ includes at least one edge of the minimum cut $D(S,T)$. Remember that the edges from $D(S,T)$ of the game Γ have the same latency functions as the edges E' of the parallel link game Γ'. Therefore, $\mathsf{OPT}(\Gamma') \leq \mathsf{OPT}(\Gamma)$ and thus $\frac{\mathsf{SC}(\mathbf{x},\Gamma)}{\mathsf{OPT}(\Gamma)} \leq \frac{\mathsf{SC}(\hat{\mathbf{x}},\Gamma')}{\mathsf{OPT}(\Gamma')}$. ■

Recall that in the case of parallel links bottleneck games with splittable traffic do not differ from Wardrop games and hence the prices of stability and anarchy coincide. This, together with Theorem 9.5 implies that the price of stability for bottleneck games with splittable traffic on arbitrary graphs corresponds to the price of stability (or anarchy) for Wardrop games on parallel links. Consequently, the results by Roughgarden [82] on the price of anarchy for Wardrop games lead to the following corollary.

Corollary 9.1: *Let $\Gamma = (G, v_o, v_d, (f_e)_{e \in E}, r)$ be a bottleneck game with splittable traffic where all functions f_e, $e \in E$, are polynomials of degree at most d with non-negative coefficients. Then there exists a Wardrop equilibrium $\mathbf{x}$ where*

$$\frac{\mathsf{SC}(\mathbf{x}, \Gamma)}{\mathsf{OPT}(\Gamma)} \leq \frac{(d+1) \cdot \sqrt[d]{d+1}}{(d+1) \cdot \sqrt[d]{d+1} - d}.$$

For readers who are familiar with the anarchy value defined in [82] we would like to mention that it is possible to draw a more general conclusion that was also given by Cole et al. [22, Theorem 4.6]: If the anarchy value $\alpha(\mathcal{F})$ exists for a set of functions $\mathcal{F}$ this value $\alpha(\mathcal{F})$ is an upper bound on the price of stability for bottleneck games with splittable traffic and latency functions from $\mathcal{F}$ on general graphs, i.e., $\mathsf{PoS}(\mathscr{G}(\mathcal{F})) \leq \alpha(\mathcal{F})$. Under some moderate assumptions being made on $\mathcal{F}$ even equality holds (see [84, Lemma 3.4.4]).

This result, however, cannot be used to prove our result on the price of stability for bottleneck games with splittable traffic and M/M/1 latency functions that we will give in Section 9.3, since we will include other game properties in order to get a meaningful result. Therefore, Theorem 9.5 is essential for the generalization from parallel links to arbitrary graphs in the M/M/1 case.

9.3 Games with M/M/1 Latency Functions

In the rest of this chapter, we will focus on bottleneck games with splittable traffic and M/M/1 latency functions. We observed earlier in Lemma 9.1 that there is such a game with infinite price of anarchy. This justifies looking at the price of stability instead. Unfortunately, also $\mathsf{PoS}(\mathscr{G}(\mathcal{M})) = \mathsf{PoS}(\mathscr{P}(\mathcal{M})) = \infty$, which will be a trivial consequence of Theorem 9.8. Hence, we need to consider other game properties, too, in order to get a meaningful result for the price of stability.

To achieve this goal we will establish the social cost of equilibria and optimum solutions for the parallel link case in Section 9.3.1 and then in Section 9.3.2 derive the exact value for $\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r))$ and then argue that it is the same as $\mathsf{PoS}(\mathscr{G}(\mathcal{M}_{\geq c}, m, r))$. By our notation, m is meant here to denote the maximum number of edges, c the minimum edge capacity, and r the maximum amount of traffic.

9.3.1 Social Cost of Equilibria and Optimum Solutions

For our later proofs on the price of stability we need some insight into the social cost of Wardrop equilibria and optimum solutions. Thus we now give the exact social cost of Wardrop equilibria for the case of parallel links.

Theorem 9.6 (Mazalov, Monien, Schoppmann, Tiemann [68]): *Let $\Gamma = (G, v_o, v_d, (c_e)_{e\in E}, r)$ be a bottleneck game with splittable traffic and M/M/1 latency functions on m parallel links where $r < C$, and let $\mathbf{x}$ be a Wardrop equilibrium. Furthermore, let $s = |\{i \in [m] \mid x_i > 0\}|$ denote the number of links used in $\mathbf{x}$. Then*

$$s = \max\left\{i \in [m] \mid r + i \cdot c_i > C^{\leq i}\right\} \quad \text{and} \quad \mathsf{SC}(\mathbf{x}, \Gamma) = \frac{s \cdot r}{C^{\leq s} - r}.$$

We will now derive an expression that describes the social cost of an optimum solution.

Theorem 9.7 (Mazalov, Monien, Schoppmann, Tiemann [68]): *Let $\Gamma = (G, v_o, v_d, (c_e)_{e\in E}, r)$ be a bottleneck game with splittable traffic and M/M/1 latency functions on m parallel links where $r < C$, and let $\mathbf{x}$ be a strategy profile with optimal social cost. Furthermore, let $t = |\{i \in [m] \mid x_i > 0\}|$ denote the number of links used in $\mathbf{x}$. Then*

$$t = \max\left\{i \in [m] \,\middle|\, r + \sqrt{c_i} \cdot \sum_{k=1}^{i} \sqrt{c_k} > C^{\leq i}\right\} \text{ and } \mathsf{OPT}(\Gamma) = \frac{\left(\sum_{i=1}^{t} \sqrt{c_i}\right)^2}{C^{\leq t} - r} - t.$$

9.3.2 Price of Stability

Combining our knowledge about the social cost of Wardrop equilibria and optimum solutions, we will now give the exact price of stability for games with a minimum edge capacity c on m parallel links routing a traffic of r.

Theorem 9.8 (Mazalov, Monien, Schoppmann, Tiemann [68]): *For bottleneck games $\Gamma = (G, v_o, v_d, (c_e)_{e\in E}, r)$ with splittable traffic at most $r > 0$ and M/M/1 latency functions on no more than $m \in \mathbb{N}$ parallel links each with a minimum capacity of at least $c > 0$ the price of stability is exactly*

$$\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r)) \quad = \quad \frac{m \cdot \frac{r}{c}}{\frac{r}{c} + 2 \cdot (m-1) \cdot \left(\sqrt{\frac{r}{c} + 1} - 1\right)}.$$

We conclude this section by the following corollary which is a direct consequence of the preceding result together with Theorem 9.5.

Corollary 9.2: *The price of stability for general bottleneck games with splittable traffic at most $r > 0$ and M/M/1 latency functions on a graph with no more than $m \in \mathbb{N}$ edges each with a minimum capacity of at least $c > 0$ is the same as in the parallel links case, i.e.,*

$$\mathsf{PoS}(\mathscr{G}(\mathcal{M}_{\geq c}, m, r)) \quad = \quad \mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r)).$$

9.4 Conclusion and Directions for Further Research

In this chapter, we studied bottleneck games with splittable traffic. Our study revealed that the social cost of Wardrop equilibria is unique for games on series parallel graphs. On any graph whose subgraph induced by all simple origin-destination paths is not series parallel, however, there exists a game having equilibria with different social cost. We furthermore established that the price of stability for bottleneck games with splittable traffic is independent of the network topology, in a similar way as the price of anarchy for Wardrop games. This result also holds for our new formula describing the exact price of stability for games with M/M/1 latency functions.

Note that the bottleneck games with splittable traffic that we studied in this chapter only have *one* origin destination pair. Up to now we do not know whether it is possible to extend our results to a multi-commodity setting with *multiple* origin destination pairs.

Bibliography

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science*, pages 613–622, 2006.

[2] H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash Equilibria in Player-Specific and Weighted Congestion Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 4286, Springer Verlag, pages 50–61, 2006.

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

[4] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 3884, Springer Verlag, pages 218–229, 2006.

[5] C. Aliprantis and K. Border. *Infinite Dimensional Analysis*. Springer Verlag, 1994.

[6] E. Altman, T. Boulogne, R. El Azouzi, T. Jiménez, and L. Wynter. A Survey on Networking Games in Telecommunications. *Computers & Operations Research*, 33:286–311, 2006.

[7] V. Anantharam. On the Nash Dynamics of Congestion Games with Player-Specific Utility. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, pages 4673 –4678, 2004.

[8] E. Anshelevich, A. Dasgupta, J. Kleinberg, T. Roughgarden, É. Tardos, and T. Wexler. The Price of Stability for Network Design with Fair Cost Allocation. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science*, pages 295–304, 2004.

[9] B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 57–66, 2005.

[10] B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. In *Proceedings of the 1st International Workshop on Approximation and Online Algorithms*, Lecture Notes in Computer Science, Vol. 2909, Springer Verlag, pages 41–52, 2003.

[11] R. Banner and A. Orda. Bottleneck Routing Games in Communication Networks. In *Proceedings of the 25th IEEE International Conference on Computer Communications*, 2006.

[12] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[13] R. Beier, A. Czumaj, P. Krysta, and B. Vöcking. Computing Equilibria for Congestion Games with (Im)perfect Information. *ACM Transactions on Algorithms*, 2(4):679–706, 2006.

[14] P. Berenbrink, L. A. Goldberg, P. W. Goldberg, and R. Martin. Utilitarian Resource Assignment. *Journal of Discrete Algorithms*, 4(4):567–587, 2006.

[15] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Second Edition. Springer Verlag, 1980.

[16] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1987.

[17] Y. Bleischwitz, B. Monien, F. Schoppmann, and K. Tiemann. The Power of Two Prices: Beyond Cross-Monotonicity. In *Proceedings of the 32nd International Symposium on Mathematical Foundations of Computer Science*, to appear 2007.

[18] C. Busch and M. Magdon-Ismail. Atomic Routing Games on Maximum Congestion. In *Proceedings of the 2nd International Conference on Algorithmic Aspects in Information and Management*, pages 79–91, 2006.

[19] I. Caragiannis, C. Galdi, and C. Kaklamanis. Network Load Games. In *Proceedings of the 16th Annual International Symposium on Algorithms and Computation*, Lecture Notes in Computer Science, Vol. 3827, Springer Verlag, pages 809–818, 2005.

[20] S. Catoni and S. Pallottino. Traffic Equilibrium Paradoxes. *Transportation Science*, 25(3):240–244, 1991.

[21] G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 67–73, 2005.

[22] R. Cole, Y. Dodis, and T. Roughgarden. Bottleneck Links, Variable Demand, and the Tragedy of the Commons. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 668–677, 2006.

[23] R. Comminetti, J. R. Correa, and N. E. Stier-Moses. Network Games with Atomic Players. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming, Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 525–536, 2006.

[24] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Computational Complexity, Fairness, and the Price of Anarchy of the Maximum Latency Problem. In *Proceedings of the 10th Conference on Integer Programming and Combinatorial Optimization*, Lecture Notes in Computer Science, Vol. 3064, Springer Verlag, pages 59–73, 2004.

[25] A. Czumaj, P. Krysta, and B. Vöcking. Selfish Traffic Allocation for Server Farms. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 287–296, 2002.

[26] A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. *ACM Transactions on Algorithms*, 3(1):Article No. 4, 2007.

[27] R. Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10:303–318, 1965.

[28] J. Dunkel and A. Schulz. On the Complexity of Pure-Strategy Nash Equilibria in Congestion and Local-Effect Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 4286, Springer Verlag, pages 62–73, 2006.

[29] B. Ellickson. *Competitive Equilibrium*. Cambridge University Press, 1993.

[30] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 604–612, 2004.

[31] G. Facchini, F. van Megan, P. Borm, and S. Tijs. Congestion Models and Weighted Bayesian Potential Games. *Theory and Decision*, 42:193–206, 1997.

[32] T. Fahle and K. Tiemann. An Improved Branch-and-Bound Algorithm for the Test Cover Problem. In *Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms*, Lecture Notes in Computer Science, Vol. 3505, Springer Verlag, pages 89–100, 2005.

[33] T. Fahle and K. Tiemann. A Faster Branch-and-Bound Algorithm for the Test Cover Problem Based on Set Covering Techniques. *ACM Journal of Experimental Algorithmics*, 11: Article 2.2, 2006.

[34] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the Coordination Ratio for a Selfish Routing Game. In *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 2719, Springer Verlag, pages 514–526, 2003.

[35] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish Routing in Non-Cooperative Networks: A Survey. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 21–45, 2003.

[36] S. Fischer, H. Räcke, and B. Vöcking. Fast Convergence to Wardrop Equilibria by Adaptive Sampling Methods. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 653–662, 2006.

[37] S. Fischer and B. Vöcking. On the Structure and Complexity of Worst-Case Equilibria. In *Proceedings of the 1st International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 3828, Springer Verlag, pages 151–160, 2005.

[38] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 2380, Springer Verlag, pages 123–134, 2002.

[39] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. In *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 593–605, 2004.

[40] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, 348(2–3):226–239, 2005.

[41] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1992.

[42] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 613–622, 2004.

[43] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 645–657, 2004.

[44] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Structure and Complexity of Extreme Nash Equilibria. *Theoretical Computer Science*, 343:133–157, 2005.

[45] M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. In *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3580, Springer Verlag, pages 51–65, 2005.

[46] M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. In *Proceedings of the 17th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 203–212, 2005. Also accepted to *Theory of Computing Systems*.

[47] M. Gairing, B. Monien, and K. Tiemann. Routing (Un-) Splittable Flow in Games with Player-Specific Linear Latency Functions. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming, Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 501–512, 2006.

[48] C. Georgiou, T. Pavlides, and A. Philippou. Network Uncertainty in Selfish Routing. In *Proceedings of the 20th IEEE International Parallel & Distributed Processing Symposium, CD-ROM Proceedings*, page 105, 2006.

[49] K. R. Goodearl. *Partially Ordered Abelian Groups with Interpolation*. American Mathematical Society, 1986.

[50] R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

[51] J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.

[52] J. C. Harsanyi. Games with Randomly Disturbed Payoffs. *International Journal on Game Theory*, 21:1–23, 1973.

[53] A. Haurie and P. Marcotte. On the Relationship Between Nash-Cournot and Wardrop Equilibria. *Networks*, 15:295–308, 1985.

[54] J. Havil. *Gamma: Exploring Euler's Constant*. Princeton University Press, 2003.

[55] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How Easy is Local Search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.

[56] L. G. Khachiyan. A Polynomial Time Algorithm in Linear Programming. *Soviet Mathematics Doklady*, 20(1):191–194, 1979.

[57] L. Kleinrock. *Queueing Systems, Volume 1: Theory*. Wiley-Interscience, 1975.

[58] Y. A. Korilis, A. A. Lazar, and A. Orda. Capacity Allocation Under Noncooperative Routing. *IEEE Transactions on Automatic Control*, 42(3):309–325, 1997.

[59] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.

[60] E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 1563, Springer Verlag, pages 404–413, 1999.

[61] F. T. Leighton, B. M. Maggs, and S. B. Rao. Packet Routing and Job-Scheduling in O(Congestion + Dilation) steps. *Combinatorica*, 14(2):167–186, 1994.

[62] L. Libman and A. Orda. Atomic Resource Sharing in Noncooperative Networks. *Telecommunication Systems*, 17(4):385–409, 2001.

[63] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 2996, Springer Verlag, pages 547–558, 2004.

[64] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash Equilibrium? In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 551–561, 2003.

[65] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[66] M. Mavronicolas, I. Milchtaich, B. Monien, and K. Tiemann. Congestion Games with Player-Specific Constants. In *Proceedings of the 32nd International Symposium on Mathematical Foundations of Computer Science*, to appear 2007.

[67] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 510–519, 2001.

[68] V. Mazalov, B. Monien, F. Schoppmann, and K. Tiemann. Wardrop Equilibria and Price of Stability for Bottleneck Games with Splittable Traffic. In *Proceedings of the 2nd International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 4286, Springer Verlag, pages 331–342, 2006.

[69] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

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

[71] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1997.

[72] J. F. Nash. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences of the United States of America*, 36:48–49, 1950.

[73] J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.

[74] A. Orda, R. Rom, and N. Shimkin. Competitive Routing in Multiuser Communication Networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

[75] M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[76] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.

[77] A. C. Pigou. *The Economics of Welfare.* Macmillan and Company, 1920.

[78] K. P. Rath. A direct proof of the existence of pure strategy equilibria in games with a continuum of players. *Economic Theory*, 2(3):427–433, 1992.

[79] T. G. Robertazzi. *Computer Networks and Systems: Queueing Theory and Performance Evaluation*. Springer-Verlag, 2000.

[80] R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[81] R. W. Rosenthal. The Network Equilibrium Problem in Integers. *Networks*, 3:53–59, 1973.

[82] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[83] T. Roughgarden. The Maximum Latency of Selfish Routing. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 980–981, 2004.

[84] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[85] T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.

[86] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.

[87] M. Schwarz. *Telecommunication Networks: Protocols, Modeling and Analysis*. Addison-Wesley, 1987.

[88] I. Szalkai and D. Velleman. Versatile Coins. *American Mathematical Monthly*, 100(1):26–33, 1993.

[89] J. Valdes. *Parsing Flowcharts and Series-Parallel Graphs*. PhD Thesis, Stanford University, 1978.

[90] B. Vöcking. Congestion Games: Optimization in Competition. In *Proceedings of the 2nd Algorithms and Complexity in Durham Workshop*, pages 9–20, 2006.

[91] M. Voorneveld, P. Borm, F. van Megan, S. Tijs, and G. Facchini. Congestion Games and Potentials Reconsidered. *International Game Theory Review*, 1:283–299, 1999.

[92] J. G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proceedings of the Institute of Civil Engineers, Part II,* Vol. 1, pages 325–378, 1952.