Selfish Routing in Networks

Dissertation

von

Martin Gairing

Schriftliche Arbeit zur Erlangung des Grades
Doktor der Naturwissenschaften
an der Fakultät für Elektrotechnik, Informatik und Mathematik
der Universität Paderborn

Paderborn, 10. Oktober 2006

To My Family

Acknowledgements

I have been very lucky to work together and learn from a number of people during the last years. First and foremost, I'm deeply indebted to my adviser, Burkhard Monien, for always offering invaluable help and encouragement. His advice, knowledge and support have been the key for achieving my goals.

Moreover, I would like to express my thanks to Thomas Lücking, Marios Mavronicolas and Karsten Tiemann for the great collaboration. Thank you to all my other co-authors on the topic of selfish routing; those are: Sebastian Aland, Dominic Dumrauf, Robert Elsässer, Rainer Feldmann, Manuel Rode, Florian Schoppmann, and Paul Spirakis. I am indebted to Ulf-Peter Schroeder for the careful reading of a preliminary version of this dissertation and for numerous suggestions for improvement. I thank all my colleagues for many discussions on scientific and non-scientific topics, for the great atmosphere, and for technical and administrative support. To name those who I did not mention yet: Ulrich Ahlers, Bernard Bauer, Yvonne Bleischwitz, Torsten Fahle, Sven Grothklags, Sigrid Gundelach, Georg Kliewer, Michael Laska, Ulf Lorenz, Henning Meyerhenke, Tomáš Plachetka, Marion Rohloff, Thomas Sauerwald, Stefan Schamberger, Norbert Sensen, Meinolf Sellmann, Thomas Thissen, Karsten Tiemann, Tobias Tscheuschner, Andreas Woclaw, and Oleksiy Znamenshchykov.

Before coming to Paderborn, I benefited from a number of mentors at Clemson University. I am especially grateful to Steven Hedetniemi for his continuous support and for introducing me to the joy of doing research. Together with David Jacobs, Renu Laskar, and Pradip Srimani, he provided support for my masters degree and encouraged me to start a PhD project.

Finally I want to thank my family and my friends for their love, support and encouragement. In particular, I am indebted to my mother and grandmother for encouraging a love of learning at an early age. Last, but not least, my love goes to my wife Ulrike and my sons Noah and Simon.

Paderborn, October 2006 *Martin Gairing*

Contents

1 Introduction 1
- 1.1 Motivation Framework 1
- 1.2 Contribution 2
  - 1.2.1 Routing Games on Parallel Links 2
  - 1.2.2 Weighted Congestion Games 2
  - 1.2.3 Selfish Routing with Incomplete Information 3
- 1.3 Related Models for Selfish Routing 3
- 1.4 Publications 4
- 1.5 Organization 4

2 Preliminaries 5
- 2.1 Notation 5
- 2.2 Gamma Function 5
- 2.3 Falling Factorials, Stirling Numbers and Bell Numbers 6
- 2.4 Binomial Cost Function 6

3 Models 9
- 3.1 Routing Games on Parallel Links 9
  - 3.1.1 Instance 9
  - 3.1.2 Strategies and Strategy Profiles 10
  - 3.1.3 Load and Latency 10
  - 3.1.4 Private Cost 10
  - 3.1.5 Social Cost Measures 11
  - 3.1.6 Nash Equilibria 13
  - 3.1.7 Price of Anarchy 13
  - 3.1.8 Selfish Steps and Nashification 14
- 3.2 Weighted Congestion Games 14
  - 3.2.1 Instance 14
  - 3.2.2 Strategies and Strategy Profiles 14
  - 3.2.3 Private Cost 15
  - 3.2.4 Nash Equilibria 15
  - 3.2.5 Social Cost 15

3.2.6 Price of Anarchy ... 16
3.2.7 Network Congestion Games ... 16
3.3 Bayesian Routing Games ... 16
3.3.1 Instance ... 16
3.3.2 Strategies and Strategy Profiles ... 17
3.3.3 Private Cost ... 18
3.3.4 Bayesian Nash Equilibria ... 19
3.3.5 Social Cost and Price of Anarchy ... 20
3.3.6 Weighted Bayesian Congestion Games ... 21

4 Selfish Routing on Parallel Links ... 23
4.1 Introduction ... 23
4.1.1 Summary of Results ... 23
4.1.2 Related Work ... 25
4.1.3 Organization ... 27
4.2 Identical Links ... 27
4.2.1 Sequences of Selfish Steps ... 27
4.2.2 Nashification ... 30
4.3 Related Links ... 32
4.3.1 Nashification ... 32
4.3.2 Price of Anarchy ... 34
4.4 Restricted Strategy Sets ... 47
4.4.1 Computation of Pure Nash Equilibria ... 47
4.4.2 Price of Anarchy ... 56
4.5 Polynomial Social Cost ... 66
4.5.1 Identical Players ... 68
4.6 Conclusion and Discussion ... 81

5 Weighted Congestion Games ... 83
5.1 Introduction ... 83
5.1.1 Summary of Results ... 83
5.1.2 Related Work ... 85
5.1.3 Organization ... 86
5.2 Price of Anarchy for Unweighted Congestion Games ... 86
5.2.1 Upper Bound ... 86
5.2.2 Lower Bound ... 93
5.3 Price of Anarchy for Weighted Congestion Games ... 95
5.3.1 Upper Bound ... 95
5.3.2 Lower Bound ... 97
5.4 Conclusion and Discussion ... 101

6 Bayesian Routing Games ... 103
6.1 Introduction ... 103
6.1.1 Summary of Results ... 104
6.1.2 Related Work ... 106

6.1.3 Organization . . . 107
6.2 Pure Bayesian Nash Equilibria . . . 107
6.2.1 Existence . . . 107
6.2.2 Computation . . . 112
6.3 Properties of Fully Mixed Bayesian Nash Equilibria . . . 114
6.4 Social Cost and Price of Anarchy . . . 120
6.4.1 Makespan Social Cost . . . 121
6.4.2 Social Cost as Sum of Private Costs . . . 129
6.4.3 Social Cost as Maximum of Private Costs . . . 131
6.5 Conclusion and Discussion . . . 134

References . . . 135

1

Introduction

1.1 Motivation Framework

Large-scale traffic and communication networks, like e.g. the Internet, telephone networks, or road traffic systems often lack a central regulation for several reasons: The size of the network may be too large, the networks may be dynamically evolving over time, or the users of the network may be free to act according to their private interests, without regard to the overall performance of the system. Besides the lack of central regulation even cooperation among the users may be impossible due to the fact that the users may not even know each other. Networks with non-cooperative users have already been studied in the early 1950's in the context of road traffic systems [7, 97]. Nowadays, modern computer artifacts, like e.g. the Internet, are modeled as communication networks with non-cooperative users. For such communication networks, combining ideas from game theory and computer science has become increasingly important [29, 63, 82, 83, 86].

An environment, which lacks a central control unit due to its size or operational mode, can be modeled as a non-cooperative game [85]. Here, the users are assumed to be selfish players that selfishly choose their private strategies, which in our environment correspond to paths (or probability distributions over the paths) from their sources to their destinations. When routing their traffic according to the strategies chosen, the players will experience an expected latency caused by the traffic of all players sharing edges. Each player tries to minimize its private cost, expressed in terms of its expected latency. This often contradicts the goal of optimizing the social cost which measures the global performance of the whole network. The degradation of the global performance due to the selfish behavior of its players is often termed as price of anarchy [86] or coordination ratio [67]. The theory of Nash equilibria [80, 81] provides us with an important concept for environments of this kind: A Nash equilibrium is a state of the system in which no player can decrease its private cost by unilaterally changing its strategy. It has been shown by Nash that a Nash equilibrium exists under fairly broad circumstances.

The concept of Nash equilibria has become an important mathematical tool in analyzing the behavior of selfish players in non-cooperative systems [86]. Many algorithms have been developed to compute a Nash equilibrium in a general game (see [76] for an overview). Although the theorem of Nash [80, 81] guarantees the existence of a Nash equilibrium, the complexity of computing a Nash equilibrium was open for a long time, even for 2-player games. Only recently, Chen and Deng [15] settled the complexity of computing a Nash equilibrium for 2-player games.

1.2 Contribution

In this work, we study different models for selfish routing in non-cooperative networks. Our models differ in the structure of the underlying network and the information accessible to the players. We now give a high-level description for the models considered in this thesis and for our contributions.

1.2.1 Routing Games on Parallel Links

In a *routing game on parallel links* a set of n players wishes to assign their *traffic* $w_1, \ldots, w_n$ to one of m parallel links connecting a single source node to a single destination node. Each link has a certain *capacity*, that represents the rate at which the link processes traffic. So the *latency* for a link is the total traffic through this link divided by its capacity. A *pure strategy* for a player is some specific link, while a *mixed strategy* is a probability distribution over its pure strategies. A *pure (resp. mixed) strategy profile* specifies a pure (resp. mixed) strategy for each player. Each player chooses a strategy in order to minimize its *private cost*, which is defined as its *expected latency*. A strategy profile is a *Nash equilibrium* if no player can decrease its private cost by unilaterally changing its strategy.

Associated with a strategy profile, there is also a global objective function, called *social cost*. For routing games on parallel links we consider two different social cost measures: *makespan social cost* and *polynomial social cost*. On the one hand, makespan social cost is defined as the expected maximum latency on a link [67]. On the other hand, polynomial social cost is defined as the sum of a certain polynomial, evaluated at the incurred link loads [42]. The maximum ratio between the maximum social cost of a Nash equilibrium and the minimum social cost of a pure strategy profile is called *price of anarchy*.

In this dissertation, we present results concerning the computational complexity of pure Nash equilibria. Furthermore, we prove a multitude of results that are related to the price of anarchy in various sub-models.

1.2.2 Weighted Congestion Games

The class of *congestion games* has been introduced by Rosenthal [88]. In a congestion game there is a set of resources and the strategy set of each player is a subset of the power set of these resources. The *latency* on a resource is determined by a *latency function* in the number of players sharing this resource. Each player aims to minimize its *private cost*, which is defined as the sum of the latencies of its chosen resources. Milchtaich [77] considered *weighted congestion games* as an extension to congestion games in which the players have weights and thus different influence on the latency of the resources. Weighted congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem. A typical resource sharing problem is that of routing. In a routing game the strategy sets of the players correspond to paths in a network. Routing games where the demand of the players cannot be split among multiple paths are also called *(weighted) network congestion games*.

For weighted congestion games we use the *total latency* [92] as our social cost measure. For the case of network congestion games, the total latency is a measure for the weighted total travel time of the players. Given this social cost measure, the price of anarchy is defined as before.

In this dissertation, we show exact values for the price of anarchy of weighted and unweighted congestion games with polynomial latency functions. The given values also hold for weighted and unweighted network congestion games.

1.2.3 Selfish Routing with Incomplete Information

In his seminal work, Harsanyi [58] introduced an elegant approach to study non-cooperative games with *incomplete information*. In our work, we use this approach to define a new selfish routing game with incomplete information that we call *Bayesian routing game*. Here, each of n selfish *players* wishes to assign its *traffic* to one of m parallel *links*. Again, the rate at which links process traffic is given by their *capacities*. However, this time players do not know each other's traffic. Following Harsanyi's approach, we introduce, for each player, a set of possible *types*. In our model, each type of a player corresponds to some traffic and the players' uncertainty about each other's traffic is described by a probability distribution over all possible *type profiles*.

In this dissertation, we prove results on the existence and computational complexity of pure Bayesian Nash equilibria, we study structural properties of a certain class of mixed Bayesian Nash equilibria, and we prove bounds on the price of anarchy for various social cost measures.

1.3 Related Models for Selfish Routing

The *Wardrop model* has already been studied in the 1950's in the context of road traffic systems by Wardrop [97] and by Beckmann, McGuire and Winsten [7]. Moreover, it was already discussed earlier by Pigou [87] in the 1920's. For a survey of the early work on this model see [8]. In the Wardrop model, traffic has to be sent through a shared network and traffic is allowed to be split into arbitrary pieces. In this environment, unregulated traffic is modeled as network flow. Wardrop [97] introduced the concept of *Wardrop equilibria* to describe user behavior in this kind of traffic networks. Given an arbitrary network with edge latency functions, Wardrop equilibria have been classified as flows with all flow paths used between a given source-destination pair having equal latency. A Wardrop equilibrium can be interpreted as a Nash equilibrium in a game with infinitely many players, each carrying an infinitesimal amount of traffic from a source to a destination.

A lot of work (see [92, Sec. 1.2] for a brief survey) on this model has been motivated by Braess's Paradox [13]. Inspired by the arisen interest in the price of anarchy, Roughgarden and Tardos re-investigated the Wardrop model [89, 92]. Other recent work on the price of anarchy in the Wardrop model and its variations include [18, 20, 24, 50, 75]. Recently, the convergence towards a Wardrop equilibrium was studied by Fischer et al. [33, 34, 35]. Many results on the Wardrop model have been collected in the book of Roughgarden [90].

Another model for selfish routing was first discussed by Orda et al. [84] and further studied by Roughgarden [91] and Comminetti et al. [19]. In this model, the number of players is finite and each player controls a non-negligible amount of flow that can be split over different paths. In contrast to the Wardrop model, each player centrally controls its flow share, seeking to minimize the total latency of its flow share.

1.4 Publications

The results presented in this thesis are published in parts as joint work in the Proceedings of the *International Colloquium on Automata, Languages, and Programming (ICALP)* [30, 45, 48], the Proceedings of the *Italian Conference on Theoretical Computer Science (ICTCS)* [46], the Proceedings of the *International Symposium on Mathematical Foundations of Computer Science (MFCS)* [31, 42], the Proceedings of the *International Symposium on Theoretical Aspects of Computer Science (STACS)* [4], the Proceedings of the *Annual ACM Symposium on Theory of Computing (STOC)* [41], the Proceedings of the *Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)* [49], *Parallel Processing Letters* [44], and *Theoretical Computer Science* [47].

1.5 Organization

After a brief description of some basic notations, definitions and technical results in Chapter 2, we formally introduce the considered models for selfish routing in Chapter 3. In Chapter 4, we study routing games on parallel links. Chapter 5 holds our results for weighted congestion games and Chapter 6 comprises our findings for Bayesian routing games.

2

Preliminaries

This section presents some basic notations, definitions and preliminary technical results which are needed throughout this thesis.

2.1 Notation

For any integer $k \geq 1$, denote $[k] = \{1, \ldots, k\}$ and $[k]_0 = \{0, \ldots, k\}$. Furthermore, for any two integers ℓ, k with $0 \leq \ell \leq k$, denote $[\ell, k] = \{\ell, \ldots, k\}$.

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 let $(\mathsf{v_{-i}}, v_i') = (v_1, \ldots, v_{i-1}, v_i', v_{i+1}, \ldots, v_n)$.

For an event E in the sample space, denote by $\mathbf{Pr}(E)$ the probability of event E happening.

For a random variable X with associated distribution $\mathbf{P}$, denote by $\mathbf{E_P}(X)$ the the expectation of X.

2.2 Gamma Function

Denote Γ the *Gamma function*, that is, for any natural number N, $\Gamma(N+1) = N!$, while for any arbitrary real number $x > 0$,

$$\Gamma(x) = \int_{t=0}^{\infty} t^x \mathsf{e}^{-t} \, \mathsf{d}t.$$

The Gamma function is invertible, both Γ and its inverse Γ^{-1} are increasing. It is well known (see e.g. Gonnet [55]) that

$$\Gamma^{-1}(N) = \frac{\log(N)}{\log\log(N)} \cdot (1 + o(1)).$$

We will use the facts that $\Gamma(x+1) = x \cdot \Gamma(x)$ for all $x > 0$ and that $\Gamma(x) \leq x$ for all $1 \leq x \leq 3$. For an introduction to the Gamma function we refer to [60].

2.3 Falling Factorials, Stirling Numbers and Bell Numbers

For any pair of integers $k \geq 1$ and $t \geq 1$, the *t'th falling factorial of k,* denoted as $k^{\underline{t}}$, is given by

$$k^{\underline{t}} = k \cdot (k-1) \cdot \ldots \cdot (k-(t-1)),$$

when $k \geq t$. Otherwise $(t \geq k+1)$, $k^{\underline{t}} = 0$.

For any pair of integers $d \geq 1$ and $t \in [d]_0$, the *Stirling number of the second kind* [94], denoted as $S(d,t)$, counts the number of partitions of a set with d elements into exactly t *blocks* (non-empty subsets). In particular, $S(d,1) = 1$. Also, for all integers $d \geq 2$, $S(d,2) = 2^{d-1} - 1$. Stirling numbers of the second kind satisfy the recurrence relation

$$S(d,t) = \sum_{q \in [t,d]} \binom{d-1}{q-1} \cdot S(q-1, t-1)$$

for all integers $d \geq 2$ and $t \in [d]$ (see, e.g., [57, Table 265, Identity (6.15)]). It is also known that for all integers $d \geq 2$ and $k \geq 1$, $k^d = \sum_{t \in [d]} S(d,t) \cdot k^{\underline{t}}$.

For any integer $d \geq 1$, the *Bell number* of order d [10], denoted as B_d, counts the number of partitions of a set with d elements into blocks. So, clearly, $B_0 = 1$ and $B_d = \sum_{t \in [d]} S(d,t)$.

2.4 Binomial Cost Function

Definition 2.1. *For any integer $r \geq 1$, consider a probability vector $\mathbf{p} = (p_1, \ldots, p_r)$. Fix a function $g(\lambda) : \mathbb{R} \to \mathbb{R}$. Then, the* binomial function $\mathsf{BF}(\mathbf{p}, g)$ *is given by*

$$\mathsf{BF}(\mathbf{p}, g) = \sum_{A \subseteq [r]} \left(\prod_{k \in A} p_k \cdot \prod_{k \notin A} (1 - p_k) \cdot g(|A|) \right).$$

Strictly speaking, Definition 2.1 defines a *functional.* If all probabilities have the same value p, then we (abuse notation to) write $\mathsf{BF}(p, r, g)$. Clearly, in this case,

$$\mathsf{BF}(p, r, g) = \sum_{k \in [r]_0} \binom{r}{k} p^k (1-p)^{r-k} g(k).$$

We show that when g is monomial, the binomial function takes a special form.

Proposition 2.2. *For each integer $d \geq 1$,*

$$\mathsf{BF}(p, r, \lambda^d) = \sum_{t \in [d]} p^t \cdot S(d,t) \cdot r^{\underline{t}}.$$

Proof. By induction on r. For the basis case, let $r = 1$. Then, $\mathsf{BF}\left(p, 1, \lambda^d\right) = \binom{1}{1}p^1 1^d = p$ and $\sum_{t\in[d]} p^t S(d,t) 1^{\underline{t}} = p^1 S(d,1) 1^{\underline{1}} = p$, so that the claim follows.

Assume inductively that for some integer $r \geq 2$, for each integer $d \geq 1$,

$$\mathsf{BF}(p, r-1, \lambda^d) = \sum_{t\in[d]} p^t \cdot S(d,t) \cdot (r-1)^{\underline{t}}\,.$$

For the induction step, we derive that

$$\begin{aligned}
&\mathsf{BF}(p, r, \lambda^d)\\
&= \sum_{k\in[r]_0} \binom{r}{k} p^k (1-p)^{r-k} k^d\\
&= \sum_{k\in[r]} \binom{r}{k} p^k (1-p)^{r-k} k^d\\
&= \sum_{k\in[r]} \frac{r}{k} \binom{r-1}{k-1} p^k (1-p)^{r-k} k^d\\
&= p \cdot r \cdot \sum_{k\in[r]} \binom{r-1}{k-1} p^{k-1} (1-p)^{r-k} k^{d-1}\\
&= p \cdot r \cdot \sum_{k\in[r-1]_0} \binom{r-1}{k} p^k (1-p)^{r-1-k} (k+1)^{d-1}\\
&= p \cdot r \cdot \sum_{k\in[r-1]_0} \binom{r-1}{k} p^k (1-p)^{r-1-k} \left(\sum_{q\in[d-1]_0} \binom{d-1}{q} k^q \right)\\
&= p \cdot r \cdot \sum_{q\in[d-1]_0} \binom{d-1}{q} \left(\sum_{k\in[r-1]_0} \binom{r-1}{k} p^k (1-p)^{r-1-k} k^q \right)\\
&= p \cdot r \cdot \sum_{q\in[d-1]_0} \binom{d-1}{q} \mathsf{BF}(p, r-1, \lambda^q)\\
&= p \cdot r \cdot \binom{d-1}{0} \mathsf{BF}(p, r-1, 1) + p \cdot r \cdot \sum_{q\in[d-1]} \binom{d-1}{q} \mathsf{BF}(p, r-1, \lambda^q)\\
&= p \cdot r + p \cdot r \cdot \sum_{q\in[d-1]} \binom{d-1}{q} \mathsf{BF}(p, r-1, \lambda^q)\\
&= p \cdot r + p \cdot r \cdot \sum_{q\in[d-1]} \binom{d-1}{q} \left(\sum_{t\in[q]} p^t \cdot S(q,t) \cdot (r-1)^{\underline{t}} \right)\\
&= p \cdot r + \sum_{q\in[d-1]} \binom{d-1}{q} \left(\sum_{t\in[q]} p^{t+1} \cdot S(q,t) \cdot r^{\underline{t+1}} \right)
\end{aligned}$$

$$
\begin{aligned}
&= p \cdot r + \sum_{t \in [d-1]} p^{t+1} \cdot r^{\underline{t+1}} \cdot \left(\sum_{q \in [t, d-1]} \binom{d-1}{q} \cdot S(q, t) \right) \\
&= p \cdot r + \sum_{t \in [2,d]} p^{t} \cdot r^{\underline{t}} \cdot \left(\sum_{q \in [t, d]} \binom{d-1}{q-1} \cdot S(q-1, t-1) \right) \\
&= p \cdot r + \sum_{t \in [2,d]} p^{t} \cdot r^{\underline{t}} \cdot S(d, t) \\
&= \sum_{t \in [d]} p^{t} \cdot r^{\underline{t}} \cdot S(d, t) \,,
\end{aligned}
$$

as needed. ■

Proposition 2.2 implies that for a *constant* probability vector and a monomial function, the binomial function is a combinatorial sum of Stirling numbers of the second kind.

It is known [45, Lemma 3] that in case g is *convex,* the binomial function does not decrease when replacing all probabilities in the probability vector $\mathbf{p}$ by the average probability $\widetilde{p} = \frac{\sum_{i \in [r]} p_i}{r}$.

Lemma 2.3 (Gairing et al. [45]). *For a convex function* g, $\mathsf{BF}(\mathbf{p}, g) \leq \mathsf{BF}(\widetilde{p}, r, g)$.

3

Models

We now introduce the considered models for selfish routing. First, we define *routing games on parallel links* in Section 3.1. Afterwards, we introduce the class of *weighted congestion games* in Section 3.2. Weighted congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem (including that of routing). In Section 3.3, we introduce *Bayesian routing games on parallel links*. In this model, players have only incomplete information about each others traffic.

Each model is introduced in a self-contained fashion. A reader, who is only interested in one of the models, might skip the other two sections here.

3.1 Routing Games on Parallel Links

3.1.1 Instance

In a *routing game on parallel links*, we have a simple network consisting of a set of m parallel links connecting a *source* node to a *destination node*. Each of n players wishes to route a particular amount of traffic along a (non-fixed) link from source to destination. Assume throughout that $n \geq 2$ and $m \geq 2$. Associated with each player i is a *strategy set* $S_i \subseteq [m]$, as the set of allowed links for player i. If $S_i = [m]$ for all players $i \in [n]$, then we have *unrestricted strategy sets*, otherwise *restricted strategy sets*. Denote w_i the *traffic* of player $i \in [n]$. Define the $n \times 1$ *traffic vector* $\mathbf{w}$ in the natural way. In the model of *identical players*, all players' traffic is equal to 1. The players' traffic may vary arbitrarily in the model of *arbitrary players*. Without loss of generality assume that $w_1 \geq \ldots \geq w_n$. Let $W = \sum_{i\in[n]} w_i$.

Denote $c_j > 0$ the *capacity* of link $j \in [m]$, representing the rate at which the link processes traffic. So, the *latency* for traffic w through link j equals w/c_j. Define the $m \times 1$ *capacity vector* $\mathbf{c}$ in the natural way. In the model of *identical links*, all link capacities are equal to 1. Link capacities may vary arbitrarily in the model of *related links*. Let $C = \sum_{j\in[m]} c_j$.

An *instance* is described by a tuple $\langle \mathbf{w}, \mathbf{c} \rangle$. In case of identical players, we replace $\mathbf{w}$ by n and in case of identical links, we replace $\mathbf{c}$ by m.

3.1.2 Strategies and Strategy Profiles

A *pure strategy* for player i is some specific link ℓ_i from its strategy set S_i. A *mixed strategy* for player $i \in [n]$ is a probability distribution over pure strategies. Thus, a mixed strategy is a probability distribution over the set of links.

A *pure strategy profile* is represented by an n-tuple $\mathbf{L} = (\ell_1, \ldots, \ell_n) \in [m]^n$ while a *mixed strategy profile* is represented by an $n \times m$ *probability matrix* $\mathbf{P}$ of nm probabilities p_{ij}, $i \in [n]$ and $j \in [m]$, where p_{ij} is the probability that player i chooses link j. The *support* of player $i \in [n]$ in the mixed strategy profile $\mathbf{P}$, denoted by $\mathsf{support}_i(\mathbf{P})$, is the set of links to which player i assigns its traffic with positive probability. Thus,

$$\mathsf{support}_i(\mathbf{P}) = \{j \in [m] \mid p_{ij} > 0\}.$$

If $\mathsf{support}_i(\mathbf{F}) = [m]$, for a mixed strategy profile $\mathbf{F}$ and for all players $i \in [n]$, then we say that $\mathbf{F}$ is a *fully mixed strategy profile*. In other words, $\mathbf{F}$ is a fully mixed strategy profile, if $f_{ij} > 0$ for all players $i \in [n]$ and links $j \in [m]$.

3.1.3 Load and Latency

Fix a mixed strategy profile $\mathbf{P}$. Denote by $\delta_j(\mathbf{P})$ the *expected load* on link $j \in [m]$. Thus,

$$\delta_j(\mathbf{P}) = \sum_{i \in [n]} p_{ij} w_i.$$

In the same way, denote by $\delta_j^{-k}(\mathbf{P})$ the expected load of all players $i \in [n], i \neq k$ on link $j \in [m]$. Thus,

$$\delta_j^{-k}(\mathbf{P}) = \sum_{i \in [n], i \neq k} p_{ij} w_i.$$

Denote by $\Lambda_j(\mathbf{P})$ the *expected latency* on link $j \in [m]$. Clearly,

$$\Lambda_j(\mathbf{P}) = \frac{\delta_j(\mathbf{P})}{c_j}.$$

The *maximum expected latency* $\Lambda(\mathbf{P})$ is the maximum, over all links, of the expected latency $\Lambda_j(\mathbf{P})$ on a link $j \in [m]$, that is,

$$\Lambda(\mathbf{P}) = \max_{j \in [m]} \Lambda_j(\mathbf{P}).$$

3.1.4 Private Cost

For a pure strategy profile $\mathbf{L} = (\ell_1, \ldots, \ell_n)$, the *latency cost for player i*, denoted by $\lambda_i(\mathbf{L})$, is

$$\lambda_i(\mathbf{L}) = \frac{\sum_{k \in [n]: \ell_k = \ell_i} w_k}{c_{\ell_i}},$$

that is, the latency cost for player i is the latency of the link it chooses.

Fix now a mixed strategy profile $\mathbf{P}$. The *expected latency cost* for player $i \in [n]$ on link $j \in S_i$, denoted by $\lambda_{ij}(\mathbf{P})$, is the expectation, over all random choices of the remaining players, of the latency cost on link j given that player i assigns its traffic to link $j \in S_i$. Thus,

$$\lambda_{ij}(\mathbf{P}) = \frac{w_i + \sum_{k\in[n],k\neq i} p_{kj} w_k}{c_j} = \frac{(1-p_{ij})w_i + \delta_j(\mathbf{P})}{c_j}.$$

For each player $i \in [n]$, the *minimum expected latency cost,* denoted by $\lambda_i(\mathbf{P})$, is the minimum, over all links $j \in S_i$, of the expected latency cost for player i on link j. Thus,

$$\lambda_i(\mathbf{P}) = \min_{j\in S_i} \lambda_{ij}(\mathbf{P}) .$$

The *private cost* of player $i \in [n]$, denoted by $\mathsf{PC}_i(\mathbf{P})$, is the expected latency of player i. Thus,

$$\mathsf{PC}_i(\mathbf{P}) = \sum_{j\in[m]} p_{ij} \cdot \lambda_{ij}(\mathbf{P}).$$

Denote by $\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{P})$ the *maximum individual cost* which is the maximum, over all players, of the private costs. Thus,

$$\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{P}) = \max_{i\in[n]} \mathsf{PC}_i(\mathbf{P}).$$

3.1.5 Social Cost Measures

Associated with an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and mixed strategy profile $\mathbf{P}$ is the *social cost* [67, Section 2]. For routing games on parallel links we consider two different social cost measures.

3.1.5.1 Makespan Social Cost

In their seminal work, Koutsoupias and Papadimitriou [67] introduced the following social cost measure. Associated with an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and a mixed strategy profile $\mathbf{P}$ is the *makespan social cost*, denoted by $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{P})$, which is the expectation, over all random choices of the players, of the maximum (over all links) latency of traffic through a link. Thus,

$$\begin{aligned}\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{P}) &= \mathbf{E_P}\left(\max_{j\in[m]} \frac{\sum_{i\in[n]:\ell_i=j} w_i}{c_j}\right) \\ &= \sum_{(\ell_1,\ldots,\ell_n)\in[m]^n} \prod_{k\in[n]} p_{k\ell_k} \cdot \left(\max_{j\in[m]} \frac{\sum_{i\in[n]:\ell_i=j} w_i}{c_j}\right).\end{aligned}$$

The displayed formulas for makespan social cost refer to a pure strategy profile $\mathbf{L} = (\ell_1, \ldots, \ell_n)$ drawn according to the probability distribution induced by the mixed strategy profile $\mathbf{P}$. Note that $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{P})$ reduces to the maximum latency through a link in the case of pure strategies.

Associated with an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ is the *makespan optimum* [67, Section 2], denoted by $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$, which is the least possible maximum (over all links) latency of traffic through a link. Thus,

$$\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = \min_{\mathbf{L} \in [m]^n} \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) \ .$$

Note, that the optimum refers to a pure strategy profile. Call a pure strategy profile $\mathbf{L}$ with $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$ *optimal*.

3.1.5.2 Polynomial Social Cost

For the model of *identical links*, another social cost function was introduced by Gairing et al. [42]. Let

$$\pi_d(\lambda) = \sum_{t \in [d]_0} a_t \cdot \lambda^t$$

be a polynomial of degree $d > 0$ with non-negative coefficients. So $a_t \geq 0$ for all $t \in [d]_0$ and $a_d > 0$. Consider the model of *identical links*. Associated with an instance $\langle \mathbf{w}, m \rangle$, a *polynomial cost function* $\pi_d(\lambda)$ and a mixed strategy profile $\mathbf{P}$ is the *polynomial social cost*, denoted by $\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})$, which is the expectation of the sum, over all links, of the polynomial cost function $\pi_d(\lambda)$ evaluated at the incurred link loads. Thus, by linearity of expectation,

$$\begin{aligned}
\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) &= \mathbf{E_P}\left(\sum_{j \in [m]} \pi_d \left(\sum_{k \in [n] \,:\, \ell_k = j} w_k \right) \right) \\
&= \sum_{j \in [m]} \mathbf{E_P} \left(\pi_d \left(\sum_{k \in [n] \,:\, \ell_k = j} w_k \right) \right) \\
&= \sum_{j \in [m]} \sum_{A \subseteq [n]} \left(\prod_{i \in A} p_{ij} \right) \left(\prod_{i \notin A} (1 - p_{ij}) \right) \pi_d \left(\sum_{k \in [n] \,:\, \ell_k = j} w_k \right) .
\end{aligned}$$

The displayed formulas for polynomial social cost refer to a pure strategy profile $\mathbf{L} = (\ell_1, \ldots, \ell_n)$ drawn according to the probability distribution induced by the mixed strategy profile $\mathbf{P}$. If we restrict to the polynomial cost function $\pi_d(\lambda) = \lambda^d$, then we write $\mathsf{SC}_{\lambda^d}(\mathbf{w}, m, \mathbf{P})$. Note that

$$\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) = \sum_{0 \leq t \leq d} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P}) \ .$$

So, polynomial social cost is a linear combination (with non-negative coefficients) of *monomial social costs*.

Associated with an instance $\langle \mathbf{w}, m\rangle$ and a polynomial cost function $\pi_d(\lambda)$ is the *polynomial optimum,* denoted by $\mathsf{OPT}_{\pi_d(\lambda)}(\mathbf{w}, m)$, which is the least possible, over all pure strategy profiles, polynomial social cost. Thus,

$$\mathsf{OPT}_{\pi_d(\lambda)}(\mathbf{w}, m) = \min_{\mathbf{L}\in[m]^n} \mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{L}) \ .$$

A (pure) strategy profile $\mathbf{L}$ such that $\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{L}) = \mathsf{OPT}_{\pi_d(\lambda)}(\mathbf{w}, m)$ will be called *optimal* (for the instance $\langle \mathbf{w}, m\rangle$ and the polynomial cost function $\pi_d(\lambda)$). The *monomial optimum* is defined as the natural special case of polynomial optimum.

3.1.6 Nash Equilibria

We are interested in a special class of mixed strategy profiles called Nash equilibria [80, 81] that we describe below. Say that a player $i \in [n]$ is *satisfied* in the mixed strategy profile $\mathbf{P}$, if $\lambda_{ij}(\mathbf{P}) = \lambda_i(\mathbf{P})$ for all links $j \in \mathsf{support}_i(\mathbf{P})$, and $\lambda_{ij}(\mathbf{P}) \geq \lambda_i(\mathbf{P})$ for all links $j \in S_i \setminus \mathsf{support}_i(\mathbf{P})$. Thus, a satisfied player has no incentive to unilaterally deviate from its mixed strategy. A player $i \in [n]$ is *unsatisfied* in the mixed strategy profile $\mathbf{P}$ if i is not satisfied for in the mixed strategy profile $\mathbf{P}$.

The mixed strategy profile $\mathbf{P}$ is a *Nash equilibrium* [67, Section 2], if each player $i \in [n]$ is satisfied. In other words $\mathbf{P}$ is a Nash equilibrium, if and only if $\mathsf{PC}_i(\mathbf{P}) \leq \mathsf{PC}_i(\mathbf{P}_{-i}, l_i)$ for all players $i \in [n]$ and all links $l_i \in S_i$. Thus, each player assigns its traffic with positive probability only to links (possibly more than one of them) for which its expected latency cost is minimized. The *fully mixed Nash equilibrium* [74], denoted by $\mathbf{F}$, is a Nash equilibrium that is a fully mixed strategy profile. We will often consider fully mixed Nash equilibrium for routing games with unrestricted strategy sets on identical links. Here, the fully mixed Nash equilibrium $\mathbf{F}$ uniquely exists and has probabilities $f_{ij} = \frac{1}{m}$ for all players $i \in [n]$ and links $j \in [m]$ [74].

3.1.7 Price of Anarchy

Let $\star \in \{\mathsf{MSP}, \pi_d(\lambda)\}$. The *price of anarchy* (also known as *coordination ratio* [67, Section 2]), denoted by PoA_*, is the supremum, over all instances $\langle \mathbf{w}, \mathbf{c}\rangle$ and Nash equilibria $\mathbf{P}$, of the ratio $\frac{\mathsf{SC}_\star(\mathbf{w},\mathbf{c},\mathbf{P})}{\mathsf{OPT}_\star(\mathbf{w},\mathbf{c})}$. Thus,

$$\mathsf{PoA}_\star = \sup_{\langle \mathbf{w},\mathbf{c}\rangle, \mathbf{P}} \frac{\mathsf{SC}_\star(\mathbf{w},\mathbf{c},\mathbf{P})}{\mathsf{OPT}_\star(\mathbf{w},\mathbf{c})}.$$

Similarly, for the *pure price of anarchy*, denoted by pPoA_*, we take the supremum over all instances $\langle \mathbf{w}, \mathbf{c}\rangle$ and *pure* Nash equilibria $\mathbf{L}$. Thus,

$$\mathsf{pPoA}_\star = \sup_{\langle \mathbf{w},\mathbf{c}\rangle, \mathbf{L}} \frac{\mathsf{SC}_\star(\mathbf{w},\mathbf{c},\mathbf{L})}{\mathsf{OPT}_\star(\mathbf{w},\mathbf{c})}.$$

In the same way, the *individual price of anarchy* is the supremum, over all instances $\langle \mathbf{w}, \mathbf{c} \rangle$ and Nash equilibria $\mathbf{P}$, of the ratio

$$\frac{\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{P})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})}.$$

3.1.8 Selfish Steps and Nashification

Given a pure strategy profile $\mathbf{L} = (\ell_1, \dots, \ell_n)$, a *selfish step* of player $i \in [n]$ is a deviation to a strategy profile $(\mathbf{L}_{-i}, \ell'_i)$ where $\mathsf{PC}_i(\mathbf{L}_{-i}, \ell'_i) < \mathsf{PC}_i(\mathbf{L})$ and $\ell'_i \in S_i$. Such a selfish step is a *greedy selfish step* if there is no strategy $\ell''_i \in S_i$ for player i such that $\mathsf{PC}_i(\mathbf{L}_{-i}, \ell''_i) < \mathsf{PC}_i(\mathbf{L}_{-i}, \ell'_i)$.

For makespan social cost Fotakis et al. [37] showed that selfish steps can be used for computing a pure Nash equilibrium with non-increased social cost. We will use the term *nashification* to denote the process of converting a pure strategy profile into a pure Nash equilibrium with non-increased social cost.

3.2 Weighted Congestion Games

3.2.1 Instance

A *weighted congestion game* Γ is a tuple

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

Here, n is the number of *players* and E is the finite set of *resources*. For every player $i \in [n]$, $w_i \in \mathbb{R}^+$ is the *weight* and $S_i \subseteq 2^E$ is the *strategy set* of player i. Denote $S = S_1 \times \ldots \times S_n$ and $S_{-i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \ldots \times S_n$. For every resource $e \in E$, the *latency function* $f_e : \mathbb{R}^+ \to \mathbb{R}^+$ describes the *latency* on resource e. We consider *polynomial latency functions* with maximum degree d and non-negative coefficients, that is, for all resources $e \in E$, the latency function is of the form $f_e(x) = \sum_{j=0}^{d} a_{e,j} \cdot x^j$ with $a_{e,j} \geq 0$ for all $j \in [d]_0$.

In a (unweighted) *congestion game*, the weights of all players are equal to 1. Thus, the latency on a resource only depends on the *number* of players choosing this resource.

3.2.2 Strategies and Strategy Profiles

A *pure strategy* for player $i \in [n]$ is some specific $s_i \in S_i$ whereas a *mixed strategy* $P_i = (p(i, s_i))_{s_i \in S_i}$ is a probability distribution over S_i, where $p(i, s_i)$ denotes the probability that player i chooses the pure strategy s_i.

A *pure strategy profile* is an n-tuple $\mathsf{s} = (s_1, \dots, s_n) \in S$ whereas a *mixed strategy profile* $\mathbf{P} = (P_1, \dots, P_n)$ is represented by an n-tuple of mixed strategies. For a mixed strategy profile $\mathbf{P}$ denote by

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

the probability that the players choose the pure strategy profile $\mathsf{s} = (s_1, \dots, s_n)$.

3.2.3 Private Cost

Fix any pure strategy profile s, and denote by $\delta_e(\mathsf{s}) = \sum_{i\in[n]:e\in s_i} w_i$ the *load* on resource $e \in E$. The *private cost* of player $i \in [n]$ in a pure strategy profile s is defined by

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

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

$$\mathsf{PC}_i(\mathbf{P}) = \sum_{\mathsf{s}\in S} p(\mathsf{s}) \cdot \mathsf{PC}_i(\mathsf{s}) \,.$$

3.2.4 Nash Equilibria

We are interested in a special class of (mixed) strategy profiles called Nash equilibria [80, 81] that we describe here. Given a weighted congestion game and an associated mixed strategy profile $\mathbf{P}$, player $i \in [n]$ is *satisfied* if the player can not improve its private cost by unilaterally changing its strategy. Otherwise, player i is *unsatisfied*. The mixed strategy profile $\mathbf{P}$ is a *Nash equilibrium* if and only if all players $i \in [n]$ are satisfied, that is, $\mathsf{PC}_i(\mathbf{P}) \le \mathsf{PC}_i(\mathbf{P}_{-i}, s_i)$ for all $i \in [n]$ and $s_i \in S_i$.

Note, that if this inequality holds for all pure strategies $s_i \in S_i$ of player i, then it also holds for all mixed strategies over S_i. Depending on the type of strategy profile, we differ between *pure* and *mixed* Nash equilibria.

3.2.5 Social Cost

Associated with a weighted congestion game Γ and a mixed strategy profile $\mathbf{P}$ is the *social cost* $\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P})$ as a measure of social welfare. In particular we use the expected total latency [92], that is,

$$\begin{aligned}\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P}) &= \sum_{\mathsf{s}\in S} p(\mathsf{s}) \sum_{e\in E} \delta_e(\mathsf{s}) \cdot f_e(\delta_e(\mathsf{s})) \\ &= \sum_{\mathsf{s}\in S} p(\mathsf{s}) \sum_{i\in[n]} \sum_{e\in s_i} w_i \cdot f_e(\delta_e(\mathsf{s})) \\ &= \sum_{i\in[n]} w_i \cdot \mathsf{PC}_i(\mathbf{P}).\end{aligned}$$

The *optimum* associated with a weighted congestion game Γ is the least possible social cost, over all pure strategy profiles $\mathsf{s} \in S$. Thus,

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

3.2.6 Price of Anarchy

The *price of anarchy*, also called *coordination ratio* and denoted by $\mathsf{PoA}_{\mathsf{TL}}$, is the supremum, over all instances Γ and Nash equilibria $\mathbf{P}$, of the ratio $\frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P})}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)}$. Thus,

$$\mathsf{PoA}_{\mathsf{TL}} = \sup_{\Gamma,\mathbf{P}} \frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P})}{\mathsf{OPT}_{\mathsf{TL}}(\Gamma)}$$

3.2.7 Network Congestion Games

In a *(weighted) network congestion game* the set of resources E corresponds to edges in a graph $G = (V, E)$. For each player $i \in [n]$ we have given an origin destination pair (o_i, d_i), where $o_i, d_i \in V$. The strategy set S_i of player $i \in [n]$ is then the set of simple paths connecting its origin o_i to its destination d_i.

3.3 Bayesian Routing Games

3.3.1 Instance

A *Bayesian routing game* is a tuple $\Gamma = (n, m, \mathbf{c}, T, \mathbf{\Psi})$. Here, each of n players wishes to assign a particular amount of traffic to one of m parallel *links* connecting a source node to a destination node. Assume throughout 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 $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$. 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 *traffic* of type t, $w(t) \geq 0$. Denote $T = T_1 \times \ldots \times T_n$, the set of all possible *type profiles*. For each player $i \in [n]$, define $\tau_i = |T_i|$ as the number of types of player i. Define $\tau = \sum_{i\in[n]} \tau_i$ as the total number of types of the players. For simplicity, we assume that the traffic of all types of the players $(w(t_i))_{t_i \in T_i, i\in[n]}$ is 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 $\mathbf{\Psi} = (\Psi(t_1, \ldots, t_n))_{(t_1,\ldots,t_n)\in T}$, called *type distribution*, over the set of type profiles T. Thus, $\mathbf{\Psi}$ is a function $\mathbf{\Psi} : T \to [0, 1]$ and $\sum_{(t_1,\ldots,t_n)\in T} \Psi(t_1, \ldots, t_n) = 1$. Denote by $\Psi(i, t)$ the probability that player i is of type t. So,

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

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

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

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

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

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. Throughout we only consider instances where $\Psi(k, t) > 0$ for all players $k \in [n]$ and all types $t \in T_k$. Denote by $W(i)$ the *expected traffic* of player $i \in [n]$. Clearly,

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

Furthermore, define the *expected total traffic* 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 traffic* of player s, given that player i has type t. So,

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

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*. For such a game, we write $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T, 1)$. Here, the set T contains only one type vector t that is used with probability 1. Complete information routing games are exactly the games introduced in Section 3.1. However, in order to emphasize the connection to Bayesian routing games, we call them differently here.

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* $P_i = (p(i, \sigma_i))_{\sigma_i \in \Sigma_i}$ for player $i \in [n]$ is a probability distribution over Σ_i. Here, $p(i, \sigma_i)$ denotes the probability that player i chooses the pure strategy σ_i.

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

$$\mathsf{support}_i(\mathbf{P}) = \{j \in [m] \,|\, \exists \sigma_i \in \Sigma_i, \exists t \in T_i \text{ with } p(i,\sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

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

$$\mathsf{support}_t(\mathbf{P}) = \{j \in [m] \,|\, \exists \sigma_i \in \Sigma_i \text{ with } p(i,\sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

Note that

$$\mathsf{support}_i(\mathbf{P}) = \bigcup_{t \in T_i} \mathsf{support}_t(\mathbf{P}).$$

A *pure strategy profile* $\boldsymbol{\sigma}$ is an n-tuple $(\sigma_1, \ldots, \sigma_n) \in \Sigma$. Call $\boldsymbol{\sigma}$ *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* $\mathbf{P} = (P_1, \ldots, P_n)$ 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 $\mathsf{support}_i(\mathbf{F}) = [m]$ for all players $i \in [n]$ and $\mathsf{support}_t(\mathbf{F}) = [m]$ for all players $i \in [n]$ and types $t \in T_i$.

3.3.3 Private Cost

3.3.3.1 Pure Strategy Profiles

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

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

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

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

Denote by $\lambda_{(i,t)}^j(\boldsymbol{\sigma}, \boldsymbol{\Psi})$ the private cost of type agent (i,t) when its traffic is assigned to link $j \in [m]$. So,

$$\lambda_{(i,t)}^j(\boldsymbol{\sigma}, \boldsymbol{\Psi}) = \frac{\delta_j^{-i}(\boldsymbol{\sigma}, (\boldsymbol{\Psi}|t_i = t)) + w(t)}{c_j}.$$

Denote by $v_{(i,t)}(\boldsymbol{\sigma}, \boldsymbol{\Psi})$ 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,

$$v_{(i,t)}(\boldsymbol{\sigma}, \boldsymbol{\Psi}) = \lambda_{(i,t)}^{\sigma_i(t)}(\boldsymbol{\sigma}, \boldsymbol{\Psi}).$$

Note that $v_{(i,t)}(\boldsymbol{\sigma}, \boldsymbol{\Psi})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player i. Finally, denote by $\mathsf{PC}_i(\sigma, \boldsymbol{\Psi})$ the *private cost* of player i. Clearly,

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

3.3.3.2 Mixed Strategy Profiles

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

$$\delta_j(\mathbf{P}, \mathbf{\Psi}) = \sum_{\boldsymbol{\sigma} \in \Sigma} \prod_{i \in [n]} p(i, \sigma_i) \cdot \delta_j(\boldsymbol{\sigma}, \mathbf{\Psi})$$

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

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

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

Denote by $\lambda_{(i,t)}^j(\mathbf{P}, \mathbf{\Psi})$ the private cost of type agent (i, t) when its traffic is assigned to link $j \in [m]$. So,

$$\lambda_{(i,t)}^j(\mathbf{P}, \mathbf{\Psi}) = \frac{\delta_j^{-i}(\mathbf{P}, (\mathbf{\Psi}|t_i = t)) + w(t)}{c_j}.$$

Denote by $v_{(i,t)}(\mathbf{P}, \mathbf{\Psi})$ 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,

$$v_{(i,t)}(\mathbf{P}, \mathbf{\Psi}) = \sum_{\sigma_i \in \Sigma_i} p(i, \sigma_i) \cdot \lambda_{(i,t)}^{\sigma_i(t)}(\mathbf{P}, \mathbf{\Psi}).$$

Note that $v_{(i,t)}(\mathbf{P}, \mathbf{\Psi})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player $i \in [n]$. Finally, denote by $\mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi})$ the *private cost* of player $i \in [n]$. Clearly,

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

3.3.4 Bayesian Nash Equilibria

A strategy profile $\mathbf{P}$ is a Bayesian Nash equilibrium, if no player has an incentive to deviate from its (mixed) strategy, that is, no player can possibly decrease its private cost when other players are sticking to their strategies. Formally, the mixed strategy profile $\mathbf{P} = (P_1, \ldots, P_n)$ is a *Bayesian Nash equilibrium* if

$$\mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi}) \leq \mathsf{PC}_i(\mathbf{P}', \mathbf{\Psi})$$

for all mixed strategy profiles $\mathbf{P}' = (P_1, \ldots, P_i', \ldots, P_n)$ and for all players $i \in [n]$. Moreover, since $v_{(i,t)}(\mathbf{P}, \mathbf{\Psi})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of player i, the above condition is equivalent to

$$v_{(i,t)}(\mathbf{P}, \mathbf{\Psi}) \leq v_{(i,t)}(\mathbf{P}', \mathbf{\Psi})$$

for all mixed strategy profiles $\mathbf{P}' = (P_1, \ldots, P_i', \ldots, P_n)$ and for all players $i \in [n]$ and types $t \in T_i$. Note that $\mathbf{P}$ is a Bayesian Nash equilibrium if and only if for all players $i \in [n]$ and types $t \in T_i$,

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

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

3.3.5 Social Cost and Price of Anarchy

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

- the *makespan social cost*, which is the expectation over all player choices and type profiles, of the maximum latency on a link. So,

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

- the *sum of private costs*,

$$\mathsf{SC}_{\mathsf{SUM}}(\Gamma, \mathbf{P}) = \sum_{i\in[n]} \mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi});$$

- the *maximum of private costs*,

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

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

$$\mathsf{PoA}_* = \sup_{\Gamma,\mathbf{P}} \frac{\mathsf{SC}_*(\Gamma, \mathbf{P})}{\mathsf{OPT}_*(\Gamma)}.$$

3.3.6 Weighted Bayesian Congestion Games

A generalization of Bayesian routing games are *weighted Bayesian congestion games with linear latency functions*. In a congestion game [88], each player $i \in [n]$ can assign its traffic 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 linear cost function $g_e(x) = a_e x + b_e$. For a Bayesian congestion game, a pure strategy profile $\boldsymbol{\sigma}$ is defined by $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ with $\sigma_i : T_i \to S_i$ for all $i \in [n]$. Thus, a pure strategy of player $i \in [n]$ maps each type $t \in T_i$ to a *set* of resources, while for a Bayesian routing game a pure strategy of player $i \in [n]$ maps each type $t \in T_i$ to a *single* link.

For a pure strategy profile $\boldsymbol{\sigma}$, 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

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

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

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

4

Selfish Routing on Parallel Links

4.1 Introduction

In this chapter, we consider routing games on parallel links. Such games have been formally introduced in Section 3.1. Here, n non-cooperative players wish to route their traffic $w_1, \ldots, w_n$ through a simple network of m parallel links with capacities $c_1, \ldots, c_m$. In the model of *identical players*, all players have equal traffic. The players' traffic may be different in the model of *arbitrary players*. In the model of *identical links*, all links have equal capacity. Link capacities may vary arbitrarily in the model of *related links*.

Each player is allowed to route its traffic along links from its strategy set. If the strategy set of each player consists of all links, then we have *unrestricted strategy sets*, otherwise *restricted strategy sets*. We assume unrestricted strategy sets, if we do not explicitly state the contrary. A *pure strategy* for a player is some specific link from its strategy set, while a *mixed strategy* is a probability distribution over pure strategies. Each player utilizes a (mixed) strategy, trying to minimize its *private cost*, which is defined as its *expected latency*. A *strategy profile* specifies a strategy for each player. Such a strategy profile is a *Nash equilibrium*, if no player can improve its expected latency by unilaterally changing its strategy. Depending on the employed player strategies, we distinguish between *pure* and *mixed* Nash equilibria. We also consider *fully mixed Nash equilibria*, where each player uses each link with strictly positive probability.

Associated with a strategy profile is also a global objective function, called social cost. In this chapter we consider two different definitions of social cost. The first one, called *makespan social cost*, is defined as the *expected maximum latency on a link*. The second one, called *polynomial social cost*, is the sum (over all links) of a certain polynomial cost function evaluated at the incurred link loads. The maximum ratio between the maximum social cost of a Nash equilibrium and the minimum social cost of a pure strategy profile is called *price of anarchy*.

4.1.1 Summary of Results

We present a multitude of results for routing games on parallel links. Our results are partitioned into two groups. The first one consists of results that are concerned with the computational complexity of pure Nash equilibria. The second one comprises our findings that are related to the price of anarchy.

4.1.1.1 Computation of Pure Nash Equilibria

It is easy to see (cf. Fotakis et al. [37]) that any selfish step decreases the *lexicographical ordering* of the link latencies. This implies that any sequence of selfish steps eventually reaches a pure Nash equilibrium. However, this does not say anything about the length of such a sequence. For the model of identical links, we obtain to following results:

- The length of a sequence of selfish steps is at most $2^n - 1$, if the players always deviate to their best link (Theorem 4.2).
- It is $\mathcal{NP}$-complete to decide, whether a given pure strategy profile can be transformed into a pure Nash equilibrium within at most k selfish steps (Theorem 4.4).
- There exists an algorithm, called NashifyIdentical, that transforms a given strategy profile into a pure Nash equilibrium in $O(n \log n)$ time using at most n selfish steps (Theorem 4.5). The algorithm does not increase makespan social cost.
- Combining the PTAS of Hochbaum and Shmoys [61] for scheduling n *jobs* on m identical *machines* with NashifyIdentical yields a PTAS for computing a pure Nash equilibrium with minimum makespan social cost (Theorem 4.6).

4.1.1.2 Price of Anarchy

Makespan Social Cost

For the case that social cost is defined as the expected maximum latency on a link, we prove a comprehensive collection of bounds on the *pure* price of anarchy for the case of unrestricted and restricted strategy sets.
For the case that strategy sets are *unrestricted*, we obtain the following:

- We introduce a structural parameter ρ that specifies the relation between the largest player traffic and the link capacities. For the model of arbitrary players and related links, we use ρ to prove an upper bound on the *pure price of anarchy* (Theorem 4.10). This upper bound is tight up to an additive constant (Theorem 4.16).
- As a corollary to Theorem 4.10 we get that, for the model of arbitrary players and related links, the pure price of anarchy is upper bounded by $\Gamma^{-1}(m)$ (Corollary 4.14). This upper bound is *asymptotically* tight (Theorem 4.17).

For the case that strategy sets are *restricted*, we prove:

- For the model of identical players with restricted strategy sets and related links, the pure price of anarchy is upper bounded by $\Gamma^{-1}(n)+1$ (Theorem 4.33). This upper bound is tight up to an additive constant, if $n = m$ (Theorem 4.32).

- For the model of arbitrary players with restricted strategy sets and identical links, the pure price of anarchy is upper bounded by $\Gamma^{-1}(m)$ (Theorem 4.35). This upper bound is tight up to an additive constant (Theorem 4.32).
- For the model of arbitrary players with restricted strategy sets and related links, the pure price of anarchy is upper bounded by m (Theorem 4.38) and lower bounded by $m-1$ (Theorem 4.37).

Polynomial Social Cost

For the case that social cost is defined as the expectation of the sum (over all links) of a certain polynomial cost function of degree $d > 0$, we prove a comprehensive collection of bounds on the price of anarchy. In particular, we show:

- For the model of identical players and two identical links, the fully mixed Nash equilibrium maximizes polynomial social cost (Theorem 4.42).
- For the model of identical players and identical links, the fully mixed Nash equilibrium maximizes polynomial social cost up to the factor $\left(1+\frac{1}{n-1}\right)^d$ (Theorem 4.44).
- For the model of identical players and identical links, the price of anarchy is upper bounded by B_d; here, B_d is the *Bell number* of order d. Our analysis first shows that B_d is an upper bound on the price of anarchy, if the polynomial cost function is the d'th power (Theorem 4.48). As a corollary we get that the same upper bound also holds for general polynomial cost functions (Corollary 4.49).
- For the model of identical players and *two* identical links, the price of anarchy is upper bounded by $2^{d-2}\left(1+\left(\frac{1}{n}\right)^{d-1}\right)$, if the polynomial cost function is the d'th power (Theorem 4.50). Moreover, this upper bound is *tight* for the sub-case of two players. As a corollary we get that the same upper bound also holds for general polynomial cost functions (Corollary 4.51).

4.1.2 Related Work

Koutsoupias and Papadimitriou [67] introduced and studied a model for selfish routing on parallel links. They defined *makespan social cost* as their social cost measure and showed the first results on the *price of anarchy*.

The price of anarchy for makespan social cost, was further studied by Mavronicolas and Spirakis [74]. In this work, they also introduced and analyzed *fully mixed Nash equilibria*. In the fully mixed Nash equilibrium, each player assigns its traffic to each link with strictly positive probability.

Tight bounds on the price of anarchy for makespan social cost were given by Czumaj and Vöcking [23] and Koutsoupias et al. [66]. They showed that the price of anarchy is $\Theta(\frac{\log m}{\log\log m})$ [23, 66] for the model of identical links and $\Theta(\frac{\log m}{\log\log\log m})$ [23] for the model of related links. Also for the model of related links, but restricting to *pure* Nash equilibria, Czumaj and Vöcking [23] showed two upper bounds of $\Gamma^{-1}(m)+1 = O(\frac{\log m}{\log\log m})$ and $O(\log(\frac{c_{\max}}{c_{\min}}))$ on the *pure price of anarchy*.

Independently of our work, Awerbuch et al. [6] also studied makespan social cost for the case of restricted strategy sets. Awerbuch et al. [6] focused on the model of arbitrary players and identical links, for which they proved that the price of anarchy is $O\left(\frac{\log m}{\log \log m}\right)$ for pure Nash equilibria and $\Theta\left(\frac{\log m}{\log \log \log m}\right)$ for all (mixed) Nash equilibria. Suri et al. [95] studied a variant of the model of parallel links with restricted strategy sets where the social cost is defined as the *total latency*. For this variant, Suri et al. [95] provided some *constant* bounds on the price of anarchy. Elsässer et al. [25] studied a further restriction of the model of parallel links with restricted strategy sets, called *interaction graphs,* where all sets of allowed links for the players have size 2. The results of Elsässer et al. [25] for their model include bounds on price of anarchy for makespan social cost. In particular, Elsässer et al. [25] proved that $\Omega\left(\frac{\log m}{\log \log m}\right)$ is still a lower bound on price of anarchy for the case of identical players and identical links in the more restricted model of interaction graphs.

Gairing et al. [42] and Lücking et al. [71] studied the *pure* price of anarchy for polynomial social cost. For identical links, Gairing et al. [42] proved that the pure price of anarchy is exactly $\frac{(2^d-1)^d}{(d-1)(2^d-2)^{d-1}}(\frac{d-1}{d})^d$, if the polynomial cost function is the d'th power. For the special case of $d = 2$, this result was shown by Lücking et al. [71]; here, the pure price of anarchy is $\frac{9}{8}$.

The *fully mixed Nash equilibrium conjecture*, which states that the fully mixed Nash equilibrium has worst social cost among all Nash equilibria, was motivated by some results from Mavronicolas and Spirakis [74] and explicitly formulated by Gairing et al. [47]. The conjecture has been proved for several particular case by Fotakis et al. [37], Gairing et al. [45, 47] and Lücking et al. [71, 72]. Fischer and Vöcking [36] presented a counterexample to the fully mixed Nash equilibrium conjecture for the case of identical links and makespan social cost.

It has been first observed by Fotakis et al. [37] that for the model of related links, Graham's LPT scheduling algorithm [56] can be used to compute a pure Nash equilibrium in polynomial time. On the other hand, Fotakis et al. [37] showed that the problem of computing a pure Nash equilibrium with minimum (or maximum, respectively) makespan social cost is $\mathcal{NP}$-hard. Fotakis et al. [37] also showed that any sequence of selfish steps converges towards a pure Nash equilibrium. Even-Dar et al. [26] studied the length of such sequences under different policies to choose the deviating player, and Goldberg [54] considered the expected length of such a sequence when a random policy is applied.

Selfish routing on parallel links is closely connected to multiprocessor scheduling. Here, pure Nash equilibria and sequences of selfish steps translate to local optima and sequences of local improvements. A schedule is said to be *jump optimal* if no job on a processor with maximum load can improve by moving to another processor [93]. Obviously, the set of pure Nash equilibria is a subset of the set of jump optimal schedules. Thus, for this model the strict upper bound $2 - 2/(m + 1)$ on the ratio between best and worst makespan of jump optimal schedules [32, 93] also holds for pure Nash equilibria. Algorithms for computing a jump optimal schedule from any given schedule have been proposed in [14, 32, 93]. The fastest algorithm has been given by Schuurman and Vredeveld [93]. It always moves the job with maximum weight from a makespan processor to a processor with minimum load, using $O(n)$ moves. However, in all algorithms the resulting jump optimal schedule is not necessarily a Nash equilibrium.

Libman and Orda [68, 69], Czumaj et al. [22] and Gairing et al. [45] considered selfish routing games on parallel links with more *general latency functions*. Libman and Orda [68, 69] allow for arbitrary increasing latency functions, while Gairing et al. [45] restrict to convex (and increasing) latency functions. Czumaj et al. [22] present a thorough study for the case of general continuous non-decreasing latency functions, with emphasis on delay functions from queuing theory.

Many results for routing games on parallel links have also been collected in the surveys of Czumaj [21], Feldmann et al. [31] and Koutsoupias [65].

4.1.3 Organization

The rest of this chapter is organized as follows. Section 4.2 deals with the case of identical links, whereas the results quoted in Section 4.3 hold for related links. Section 4.4 studies the case of restricted strategy sets. Our results for polynomial social cost are presented in Section 4.5. We conclude in Section 4.6.

4.2 Identical Links

In this section, we consider routing games on identical links. Here, we are interested in the problem of computing a pure Nash equilibrium. Basically, two different approaches can be found in the literature.

The first approach is to directly compute a pure Nash equilibrium. Fotakis et al. [37] showed that the LPT algorithm, first explored by Graham [56], yields some pure Nash equilibrium.

The second approach is to convert a given pure strategy profile into a pure Nash equilibrium without increasing the social cost. This conversion process is called *nashification*. Since selfish steps do not increase makespan social cost and any sequence of selfish steps eventually reaches a pure Nash equilibrium, selfish steps seem to be suitable for nashification.

However, our results in Section 4.2.1 and Section 4.2.2 will show, that we can't use them uncoordinated. Section 4.2.1 deals with sequences of selfish steps, whereas in Section 4.2.2 we present a nashification algorithm.

4.2.1 Sequences of Selfish Steps

In this section, we establish bounds on the maximum length of sequences of greedy selfish steps. Recall, that in a greedy selfish step, the deviating player chooses its best alternative. Afterwards, we consider the problem of deciding whether a given pure strategy profile can be transformed into a pure Nash equilibrium within a given number of selfish steps.

As discussed above, performing greedy selfish steps will eventually convert any pure strategy profile into a pure Nash equilibrium. However, this may take exponential time, even for identical links, as shown in Theorem 4.2 and Theorem 4.3.

The following lemma is crucial for proving the upper bound in Theorem 4.2.

Lemma 4.1. *Consider the model of arbitrary players and identical links. Then a greedy selfish step of an unsatisfied player i with traffic w_i makes no player k with traffic $w_k \geq w_i$ unsatisfied.*

Proof. We prove a more general result in Lemma 4.7. ■

Theorem 4.2. *Consider the model of arbitrary players and identical links. Then, for any instance $\langle \mathbf{w}, m \rangle$, the length of a sequence of greedy selfish steps is at most $2^n - 1$.*

Proof. Without loss of generality assume $w_1 \geq w_2 \geq \cdots \geq w_n$. Let $1 \leq i \leq n$. We prove by induction on i that player i can make at most 2^{i-1} greedy selfish steps.

Since w_1 is the largest traffic, and because of Lemma 4.1, player 1 can make at most one greedy selfish step. This proves the claim for $i = 1$. So assume $i \geq 2$. Due to Lemma 4.1 player i can only become unsatisfied by a move of a player with larger traffic. By induction hypothesis, the number of greedy selfish steps made by players $1, \ldots, i-1$ is at most $\sum_{k=1}^{i-1} 2^{k-1} = 2^{i-1} - 1$. This shows that player i can become unsatisfied at most $2^{i-1} - 1$ times. Since after a greedy selfish step player i becomes satisfied and since player i can be unsatisfied at the beginning, player i can make at most 2^{i-1} greedy selfish steps.
Summing up over all players, the total number of greedy selfish steps is at most $\sum_{i=1}^{n} 2^{i-1} = 2^n - 1$. This completes the proof of the theorem. ■

A corresponding lower bound on the maximum length of a sequence of greedy selfish steps was independently given by Feldmann et al. [31] and Even-Dar et al. [26]. We include the latter, since it strictly dominates the former:

Theorem 4.3 (Even-Dar et al. [26]). *Consider the model of arbitrary players and identical links. Then, there exists an instance and associated pure strategy profile for which the maximum length of a sequence of greedy selfish steps is at least*

$$\frac{\left(\frac{n}{m-1}\right)^{m-1}}{2(m-1)!} .$$

Instead of the maximum length one may ask about the minimum length of a sequence of selfish steps. In particular, one may consider whether a given pure strategy profile can be transformed into a pure Nash equilibrium with at most k selfish steps. We address this question with the following decision problem:

NASHIFY

INSTANCE:	A problem instance $\langle \mathbf{w}, \mathbf{c} \rangle$, an associated pure strategy profile $\mathbf{L}$, and a positive integer k.
QUESTION:	Is there a sequence of at most k selfish steps that transforms $\mathbf{L}$ into a pure Nash equilibrium?

If k is not part of the input, then the problem is called k-NASHIFY.

In order to prove that NASHIFY is $\mathcal{NP}$-complete, we will employ a polynomial time reduction from PARTITION. PARTITION already appeared in the original list of 21 $\mathcal{NP}$-complete problems, presented by Karp [64]. In the notation of Garey and Johnson [52], PARTITION is defined as follows:

PARTITION

INSTANCE:	A finite set A of items, a size $s(a_i) \in \mathbb{N}$ for each item $a_i \in A$, $i \in [|A|]$.
QUESTION:	Is there a subset $A' \subseteq A$ such that $\sum_{a \in A'} s(a) = \sum_{a \in A \setminus A'} s(a)$?

We are now ready to establish $\mathcal{NP}$-completeness for NASHIFY.

Theorem 4.4. NASHIFY *is $\mathcal{NP}$-complete, even for the case of two identical links.*

Proof. Clearly, NASHIFY is in $\mathcal{NP}$ since it is solvable in polynomial-time by a non-deterministic algorithm. We now prove $\mathcal{NP}$-hardness by reduction from PARTITION, that is, we employ a polynomial-time transformation from PARTITION to NASHIFY. Consider any arbitrary instance of PARTITION with $k \geq 2$ items (an instance of partition with one items is a trivial *no* instance), and let $S = \sum_{a \in A} s(a)$. From this instance construct an instance of NASHIFY as follows:

- There are $n = 5k$ players with weights

$$w_i = \begin{cases} s(a_i) & \text{if } i \in [k], \\ \frac{1}{4k} & \text{if } k+1 \leq i \leq 5k. \end{cases}$$

- There are $m = 2$ identical links.
- The pure strategy profile $\mathbf{L}$ is defined as follows: All players $i \in [3k]$ are assigned to link 1 and players $3k+1, \ldots, 5k$ are assigned to link 2.

Clearly, this is a polynomial time transformation. We prove that this is a transformation from PARTITION to NASHIFY.

(1.) The instance of PARTITION is positive:
Thus, there exists a subset $A' \subseteq A$ such that $\sum_{a \in A'} s(a) = \sum_{a \in A \setminus A'} s(a)$. Since either $|A'| \leq \frac{k}{2}$ or $|A \setminus A'| \leq \frac{k}{2}$, assume, without loss of generality, that

$|A'| \leq \frac{k}{2}$. Clearly, each player $i \in [3k]$ assigned to link 1 is unsatisfied in the constructed pure strategy profile $\mathbf{L}$. Furthermore, transferring all players that correspond to an element $a \in A'$ from link 1 to link 2 (in any order) is a sequence of at most $\frac{k}{2} < k$ selfish steps. For the resulting strategy profile $\mathbf{L}'$ we have

$$\Lambda_1(\mathbf{L}') = \Lambda_2(\mathbf{L}') = \frac{S}{2} + \frac{1}{2}.$$

This implies that $\mathbf{L}'$ is a pure Nash equilibrium so that NASHIFY is positive.

(2.) The instance of NASHIFY is positive:
Thus, there exists a sequence of at most k selfish steps that transforms the pure strategy profile $\mathbf{L}$ in the constructed instance of NASHIFY to a pure Nash equilibrium $\mathbf{L}'$. Assume that in $\mathbf{L}'$ players corresponding to a subset $A' \subseteq A$ are assigned to link j_1, players corresponding to the subset $A \setminus A' \subseteq A$ are assigned to link j_2, while the sums of traffic of players with traffic $\frac{1}{4k}$ that reside in link j_1 and link j_2 are x and $1-x$, respectively. Thus, the latencies of the links are $\Lambda_{j_1}(\mathbf{L}') = \sum_{a\in A'} s(a) + x$ and $\Lambda_{j_2}(\mathbf{L}') = \sum_{a\in A\setminus A'} s(a) + 1 - x$. Without loss of generality, assume, that $\sum_{a\in A'} s(a) \geq \sum_{a\in A\setminus A'} s(a)$.
We show that this implies $\sum_{a\in A'} s(a) - \sum_{a\in A\setminus A'} s(a) = 0$. Assume otherwise $\sum_{a\in A'} s(a) - \sum_{a\in A\setminus A'} s(a) > 0$. Since the traffic of players in A is integer, this implies $\sum_{a\in A'} s(a) - \sum_{a\in A\setminus A'} s(a) \geq 1$. Since NASHIFY is positive, we made at most k selfish steps with the players having small traffic; thus, $\frac{1}{4} \leq x \leq \frac{3}{4}$. It follows that

$$\begin{aligned}\Lambda_{j_1}(\mathbf{L}') - \Lambda_{j_2}(\mathbf{L}') &= \sum_{a\in A'} s(a) + x - \sum_{a\in A\setminus A'} s(a) - 1 + x \\ &\geq 2x \\ &\geq \frac{1}{2}.\end{aligned}$$

This implies that all remaining players with traffic $\frac{1}{4k} \leq \frac{1}{8}$ on link j_1 are unsatisfied, a contradiction to the fact that NASHIFY is positive. So $\sum_{a\in A'} s(a) - \sum_{a\in A\setminus A'} s(a) = 0$ which implies that PARTITION is positive.

This completes our reduction. ■

We remark that NASHIFY is $\mathcal{NP}$-complete in the strong sense (cf. [52, Section 4.2]) if m is part of the input. Thus, there is no pseudopolynomial-time algorithm for NASHIFY (unless $\mathcal{P} = \mathcal{NP}$). In contrast, there is a natural pseudopolynomial-time algorithm for k-NASHIFY, which exhaustively searches all sequences of k selfish steps; since a selfish step involves a (unsatisfied) player and a link for a total of at most mn choices, this algorithm can be implemented to run in $\Theta((mn)^k)$ time.

4.2.2 Nashification

We provide a polynomial-time algorithm to convert any pure strategy profile into a pure Nash equilibrium with non-increased social cost. We call our algorithm NashifyIdentical. NashifyIdentical solves NASHIFY when n selfish steps are allowed. Together with the PTAS for scheduling n jobs on m identical machines [61] this yields a PTAS for computing a best pure Nash equilibrium.

NASHIFYIDENTICAL($\mathbf{L}$)
Input: A pure strategy profile $\mathbf{L}$ of n players with traffic $w_1, \ldots, w_n$.
Output: A pure strategy profile $\mathbf{L}'$ that is a Nash equilibrium.

```
Sort the players' traffic in non-increasing order so that w_1 ≥ ... ≥ w_n.
for i ← 1 to n do
    if player i is unsatisfied then
        let player i perform a greedy selfish step;
    end if
end for
return the resulting strategy profile L'
```

Fig. 4.1. The algorithm NASHIFYIDENTICAL

The algorithm NASHIFYIDENTICAL sorts the players' traffic in non-increasing order so that $w_1 \geq \ldots \geq w_n$. Then the algorithm examines the players in order of non-increasing traffic. For each player i we let player i perform a greedy selfish step, if i is unsatisfied.

Theorem 4.5. *Given an instance $\langle \mathbf{w}, m \rangle$ and an associated pure strategy profile $\mathbf{L} = \langle l_1, \ldots, l_n \rangle$, algorithm NASHIFYIDENTICAL($\mathbf{L}$) computes a pure Nash equilibrium $\mathbf{L}'$ with social cost $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L}') \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L})$ using at most n greedy selfish steps and $O(n \log n)$ time.*

Proof. Clearly, $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L}') \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L})$, since selfish steps do not increase social cost. Furthermore, after every iteration the player that changed its strategy is satisfied and stays satisfied in subsequent iterations by Lemma 4.1. Thus $\mathbf{L}'$ is a Nash equilibrium.

The running time of algorithm NASHIFYIDENTICAL is $O(n \log n)$ for sorting the n player by their traffic, $O(m \log m)$ for constructing a heap holding all link latencies in the pure strategy profile $\mathbf{L}$, and $O(\log m)$ for updating the heap in each of the n iterations of the algorithm. Thus, the total running time is $O(n \log n + m \log m + n \log m)$. The interesting case is when $m \leq n$ (since otherwise, a single player can be assigned to each link, achieving an optimal Nash equilibrium). Thus, in the interesting case, the total running time of NASHIFYIDENTICAL is $O(n \log n)$. ■

Since it is possible to compute a pure Nash equilibrium in polynomial time, one may want to go one step further and ask, whether a pure Nash equilibrium with *minimum* makespan social cost can also be computed in polynomial time. Fotakis et al. [37] showed that this problem is $\mathcal{NP}$-complete. The next logical step is to ask for an approximation algorithm to perform this task. Here, the strong connection between multiprocessor scheduling and routing on parallel links proves useful.

Since NashifyIdentical does not increase makespan social cost, we can combine any approximation algorithm for the corresponding scheduling problem with NashifyIdentical. Hochbaum and Shmoys [61] presented a PTAS for scheduling n *jobs* on m identical *machines*. Running this PTAS on an instance $\langle \mathbf{w}, m \rangle$ yields a pure strategy profile $\mathbf{L}$ such that $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L}) \leq (1+\varepsilon)\, \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, m)$. On the other hand, applying the algorithm NashifyIdentical on $\mathbf{L}$ yields a Nash equilibrium $\mathbf{L}'$ such that $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L}') \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L})$. Thus, $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, m, \mathbf{L}') \leq (1+\varepsilon)\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, m)$. It follows that:

Theorem 4.6. *There exists a* PTAS *for computing a pure Nash equilibrium with minimum makespan social cost, in the model of identical links.*

4.3 Related Links

In this section, we consider routing games on parallel *related links*. In Section 4.3.1 we show that results from Section 4.2.2 can be generalized to the model of related links. Section 4.3.2 holds our results that are related to the price of anarchy for makespan social cost.

4.3.1 Nashification

We now consider the problem of computing a pure Nash equilibrium for the model of related links. Again, Graham's LPT algorithm [56] can be used to compute such a pure Nash equilibrium directly [37]. For related links, the makespan social cost of the Nash equilibrium computed by LPT approximates the makespan social cost of a pure Nash equilibrium with minimum social cost by a factor between 1.52 and 1.67 [40].

In this section we are interested in a nashification algorithm, that is, given a pure strategy profile, we want to compute a pure Nash equilibrium with non-increased makespan social cost. Selfish steps can also be used to compute a pure Nash equilibrium, since selfish steps do not increase makespan social cost and every sequence of selfish steps eventually reaches a pure Nash equilibrium. However, it is unknown whether selfish steps can be used to implement nashification in polynomial time. Feldmann et al. [30] chose a different approach not only based on selfish steps. Their algorithm relies on the following crucial observation:

Lemma 4.7. *Consider the model of arbitrary players and related links. Then, for any pure strategy profile, a greedy selfish step of an unsatisfied player $i \in [n]$ with weight w_i from a link $j \in [m]$ to a link $k \in [m]$ with $c_j \leq c_k$ makes no satisfied player $s \in [n]$ with weight $w_s \geq w_i$ unsatisfied.*

Proof. Let $\mathbf{L}$ and $\mathbf{L}'$ be the pure strategy profiles before and after the greedy selfish step of player i. By way of contradiction assume that some player s with traffic $w_s \geq w_i$ becomes unsatisfied due to this selfish step, and let player s be assigned to link q. Since only the loads on link j and k change due to the greedy selfish step of player i we have to show that player s cannot improve by moving to link j or if $q = k$ that s does not become unsatisfied due to the arrival of player i. We proceed by case study:

- Assume first, $q \neq k$. As player s is satisfied in $\mathbf{L}$,
$$\frac{\delta_k(\mathbf{L}) + w_s}{c_k} \geq \frac{\delta_q(\mathbf{L})}{c_q}.$$
Player i improves by moving to link k, thus,
$$\frac{\delta_j(\mathbf{L})}{c_j} > \frac{\delta_k(\mathbf{L}) + w_i}{c_k} = \frac{\delta_k(\mathbf{L}')}{c_k}.$$
It follows
$$\begin{aligned}\frac{\delta_j(\mathbf{L}') + w_s}{c_j} &= \frac{\delta_j(\mathbf{L}) - w_i + w_s}{c_j}\\ &> \frac{\delta_k(\mathbf{L}) + w_i}{c_k} + \frac{w_s - w_i}{c_j}\\ &= \frac{\delta_k(\mathbf{L}) + w_s}{c_k} - \frac{w_s - w_i}{c_k} + \frac{w_s - w_i}{c_j}\\ &\geq \frac{\delta_q(\mathbf{L})}{c_q} + (w_s - w_i)\left(\frac{1}{c_j} - \frac{1}{c_k}\right)\\ &\geq \frac{\delta_q(\mathbf{L}')}{c_q}.\end{aligned}$$
The last inequality holds since $c_k \geq c_j$, $w_s \geq w_i$ and $\delta_q(\mathbf{L}') = \delta_q(\mathbf{L})$. Thus, in $\mathbf{L}'$, player s cannot improve by moving to link j.
- Now assume that $q = k$. Since player i performs a greedy selfish step from link j to link k, for all links $r \in [m]$,
$$\frac{\delta_k(\mathbf{L}) + w_i}{c_k} \leq \frac{\delta_r(\mathbf{L}) + w_i}{c_r}.$$
Because of $\frac{\delta_j(\mathbf{L}) - w_i + w_s}{c_j} \geq \frac{\delta_j(\mathbf{L})}{c_j} > \frac{\delta_k(\mathbf{L}) + w_i}{c_k}$, player s cannot improve by moving to link j. Since player i performed a greedy selfish step, for all links $r \in [m] \setminus \{j\}$, we have,
$$\begin{aligned}\frac{\delta_k(\mathbf{L}')}{c_k} &= \frac{\delta_k(\mathbf{L}) + w_i}{c_k}\\ &\leq \frac{\delta_r(\mathbf{L}) + w_i}{c_r}\\ &\leq \frac{\delta_r(\mathbf{L}) + w_s}{c_r}\\ &= \frac{\delta_r(\mathbf{L}') + w_s}{c_r}\end{aligned}$$
and therefore player s cannot improve by moving to any link $r \neq j$.

■

The algorithm of Feldmann et al. [30], call it NashifyRelated, works in two phases. In the first phase, given an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and an associated pure strategy profile $\mathbf{L}$, it fills up links with small capacities with players with small traffic as close to $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L})$ as possible (but without exceeding $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L})$), and it collects all these players in a set $\mathcal{U}$. In the second phase, the algorithm performs greedy selfish steps for unsatisfied players in $\mathcal{U}$ in non-increasing order of their traffic. Lemma 4.7 allows to show that this procedure results in a pure Nash equilibrium. Implementing the algorithm in a proper way, we get:

Theorem 4.8 (Feldmann et al. [30]). *Consider the model of arbitrary players and related links. Then for any instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and associated pure strategy profile $\mathbf{L}$, a Nash equilibrium $\mathbf{L}'$ can be computed from $\mathbf{L}$ with $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}') \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L})$ using at most $O(m^2 n)$ time.*

Hochbaum and Shmoys [62] presented a PTAS for scheduling n jobs on m related machines. Using the same arguments as in Section 4.2.2, we can combine this PTAS with NashifyRelated to get a PTAS for computing a pure Nash equilibrium with minimum makespan social cost.

Theorem 4.9 (Feldmann et al. [30]). *There exists a* PTAS *for computing a pure Nash equilibrium with minimum makespan social cost, in the model of related links.*

4.3.2 Price of Anarchy

In this section we state results on the price of anarchy and the individual price of anarchy for the case of related links and makespan social cost.

In this scenario, the first results on the price of anarchy were give by Koutsoupias and Papadimitriou [67]. For the special case of 2 links, they showed that the price of anarchy is the *golden ratio*. For the general case, Czumaj and Vöcking [23] proved that the price of anarchy is $\Theta(\frac{\log m}{\log \log \log m})$. To show this asymptotically tight upper bound, they first provided upper bounds on the maximum expected latency $\Lambda(\mathbf{P})$ on a link in a mixed Nash equilibrium $\mathbf{P}$. A Chernoff bound then gives the upper bound on the price of anarchy. The upper bounds on $\Lambda(\mathbf{P})$, given by Czumaj and Vöcking [23], depend on the number of links m and the fraction of the largest and the smallest link capacity. However, not only the capacities, but the relation between the players' traffic and the capacities determine the (individual) price of anarchy.

To take this relation into account, we introduce a structural parameter ρ. We denote

$$\mathcal{M}_1 = \{j \in [m] \mid w_1 \leq c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})\}.$$

Using $\mathcal{M}_1$, we define

$$\rho = \frac{\sum_{j \in \mathcal{M}_1} c_j}{C}. \tag{4.1}$$

In other words, ρ is the ratio between the sum of link capacities of links to which the largest traffic can be assigned causing latency at most $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ and the sum of all link capacities.

With the help of ρ we are able to prove an upper bound of $\Gamma^{-1}(\frac{1}{\rho})$ on the individual price of anarchy (Theorem 4.10). Since $\frac{w_1}{c_1} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ and $C = \sum_{j\in[m]} c_j \leq m \cdot c_1$, it follows that $\rho \geq \frac{1}{m}$. Using this, we can upper bound the maximum expected latency on a link in a mixed Nash equilibrium by $\Lambda(\mathbf{P}) \leq \Gamma^{-1}(m) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ (Corollary 4.12), which slightly improves the best known upper bound of $(\Gamma^{-1}(m)+1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ by Czumaj and Vöcking [23] and thus leads to an improvement of the upper bound on the price of anarchy. Furthermore, it follows that the individual coordination ratio is upper bounded by $\Gamma^{-1}(m)$ (Corollary 4.13). For pure Nash equilibria $\mathbf{L}$, we have $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c},\mathbf{L}) = \mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{L})$. It follows that the pure price of anarchy for makespan social cost is upper bounded by $\Gamma^{-1}(m)$ (Corollary 4.14).

We close this section with two lower bounds on the individual price of anarchy. These lower bounds show that the upper bound from Theorem 4.10 is tight up to an additive constant for *all* m (Theorem 4.16), whereas the upper bound from Corollary 4.13 is only tight for *large* m (Theorem 4.17).

Theorem 4.10. *Consider the model of arbitrary players and related links. Then for any instance* $\langle \mathbf{w},\mathbf{c}\rangle$ *and associated Nash equilibrium* $\mathbf{P}$*,*

$$\frac{\mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})} < \begin{cases} \frac{3}{2} + \sqrt{\frac{1}{\rho} - \frac{3}{4}} & \text{if } \frac{1}{3} \leq \rho \leq 1, \\ 2 + \sqrt[3]{\frac{1}{\rho} - 2} & \text{if } \frac{1}{37} \leq \rho < \frac{1}{3}, \\ \Gamma^{-1}\left(\frac{1}{\rho}\right) & \text{if } \rho < \frac{1}{37}. \end{cases}$$

Proof. Consider an arbitrary instance $\langle \mathbf{w},\mathbf{c}\rangle$ with associated mixed Nash equilibrium $\mathbf{P}$ such that

$$\mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P}) \;=\; k \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$$

for some $k \in \mathbb{R}^+$. Furthermore, let $\mathbf{L}$ be a pure strategy profile with optimum makespan social cost; thus, $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c},\mathbf{L}) = \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$. Note that there always exists such a pure strategy profile. We proceed as follows. In part (1.) and (2.) we prove a lower bound on the total amount of traffic that is necessary for the Nash equilibrium $\mathbf{P}$. In part (3.) we then use this lower bound to prove an upper bound on k for each of the three cases. We continue with the details of the formal proof.

(1.) Let j_1 be the maximum index of a link in $\mathcal{M}_1$, that is, $\mathcal{M}_1 = [j_1]$. Let $i_1 \in [n]$ be a player and let $s_1 \in \mathcal{M}_1$ be a link with $p_{i_1 s_1} > 0$ and

$$\lambda_{i_1 s_1}(\mathbf{P}) = \frac{\delta_{s_1}^{-i_1}(\mathbf{P}) + w_{i_1}}{c_{s_1}} = \mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P}).$$

Since $\mathbf{P}$ is a Nash equilibrium, we have for all links $j \in \mathcal{M}_1$,

$$\begin{aligned}\mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P}) &= \frac{\delta_{s_1}^{-i_1}(\mathbf{P}) + w_{i_1}}{c_{s_1}}\\ &\leq \frac{\delta_j^{-i_1}(\mathbf{P}) + w_{i_1}}{c_j}\\ &\leq \frac{\delta_j^{-i_1}(\mathbf{P}) + w_1}{c_j}. \end{aligned} \tag{4.2}$$

Furthermore, by definition of $\mathcal{M}_1$, we have for all links $j \in \mathcal{M}_1$,

$$\frac{w_1}{c_j} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \tag{4.3}$$

This implies that $w_1 \leq c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ for all $j \in \mathcal{M}_1$, and thus

$$\begin{aligned} k \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) &\leq \mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P})\\ &\overset{(4.2)}{\leq} \frac{\delta_j^{-i_1}(\mathbf{P}) + w_1}{c_j}\\ &\overset{(4.3)}{\leq} \frac{\delta_j^{-i_1}(\mathbf{P}) + c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})}{c_j}\\ &= \frac{\delta_j^{-i_1}(\mathbf{P})}{c_j} + \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}), \end{aligned}$$

or equivalently

$$\delta_j^{-i_1}(\mathbf{P}) \geq (k-1) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}).$$

Therefore, for all $j \in \mathcal{M}_1$,

$$\begin{aligned} \delta_j(\mathbf{P}) &= \delta_j^{-i_1}(\mathbf{P}) + p_{i_1 j} \cdot w_{i_1}\\ &\geq \delta_j^{-i_1}(\mathbf{P})\\ &\geq (k-1) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \end{aligned} \tag{4.4}$$

Summing up all expected loads $\delta_j(\mathbf{P})$ on links in $\mathcal{M}_1$, the total expected traffic of links in $\mathcal{M}_1$ is

$$\begin{aligned} \sum_{j \in \mathcal{M}_1} \delta_j(\mathbf{P}) &\overset{(4.4)}{\geq} (k-1) \cdot \sum_{j \in \mathcal{M}_1} c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\\ &\overset{(4.1)}{=} (k-1) \cdot \rho \cdot C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \end{aligned} \tag{4.5}$$

(2.) We prove an inductive claim:

Lemma 4.11. *For all integers l with $2 \leq l \leq \lceil k \rceil - 1$, there is a set of links $\mathcal{M}_l = [j_l] \setminus [j_{l-1}]$ such that* *(a) the total capacity of links in* $\mathcal{M}_l$ *is at least:*

$$\sum_{j\in\mathcal{M}_l} c_j \geq \rho \cdot C \cdot (k-2) \cdot \prod_{j=2}^{l-1} (k-j),$$

(b) the expected load on each link j *in* $\mathcal{M}_l$ *is at least:*

$$\delta_j(\mathbf{P}) \geq (k-l) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),$$

(c) the total expected load on links in $\mathcal{M}_l$ *is at least:*

$$\sum_{j\in\mathcal{M}_l} \delta_j(\mathbf{P}) \geq \rho \cdot C \cdot (k-2) \cdot \prod_{j=2}^{l} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),$$

(d) the difference between the total expected load on links in $\mathcal{M}_1 \cup \cdots \cup \mathcal{M}_l = [j_l]$ *in the mixed Nash equilibrium* $\mathbf{P}$ *and the maximum total expected load on the same links in the optimum strategy profile* $\mathbf{L}$ *is at least:*

$$\sum_{j\in[j_l]} \delta_j(\mathbf{P}) - \sum_{j\in[j_l]} \delta_j(\mathbf{L}) \geq \rho \cdot C \cdot (k-2) \cdot \prod_{j=2}^{l} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}).$$

Proof. We will first show that the claim holds for $l=2$. Let w_{i_2} be the smallest traffic of a player i_2 who chooses a link in $\mathcal{M}_1$ with positive probability, and let $s_2 \in \mathcal{M}_1$ be a link in $\mathcal{M}_1$ with $p_{i_2 s_2} > 0$. In the pure strategy profile $\mathbf{L}$ (with optimum social cost) at most

$$\begin{aligned}\sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{L}) \ &\leq \ \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \cdot \sum_{j\in\mathcal{M}_1} c_j \\ &\overset{(4.1)}{=} \rho \cdot C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \end{aligned} \tag{4.6}$$

total load can be assigned to links in $\mathcal{M}_1$. Therefore, in $\mathbf{L}$ the remaining expected load which is greater or equal to

$$\begin{aligned}\sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) - \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{L}) &\overset{(4.6)}{\geq} \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) - \rho \cdot C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\ &\overset{(4.5)}{\geq} \rho \cdot C \cdot (k-2) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \end{aligned} \tag{4.7}$$

is assigned to links not in $\mathcal{M}_1$. This implies that there exists a set of links $\mathcal{M}_2 = [j_2] \setminus [j_1]$, j_2 minimal, with total capacity at least

$$\begin{aligned}\sum_{j\in\mathcal{M}_2} c_j \ &\geq \ \frac{\sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) - \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})} \\ &\overset{(4.7)}{\geq} \rho \cdot C \cdot (k-2), \end{aligned} \tag{4.8}$$

proving (a). Moreover, since $\mathbf{P}$ is a Nash equilibrium, for all links $j \in \mathcal{M}_2$,

$$\begin{aligned}(k-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) &\overset{(4.4)}{\leq} \frac{\delta_{s_2}(\mathbf{P})}{c_{s_2}} \\ &\leq \frac{\delta_{s_2}^{-i_2}(\mathbf{P}) + w_{i_2}}{c_{s_2}} \\ &\leq \frac{\delta_{j}^{-i_2}(\mathbf{P}) + w_{i_2}}{c_j}. \end{aligned} \tag{4.9}$$

By construction of $\mathcal{M}_2$ there exists a player assigned in $\mathbf{P}$ with positive probability to a link in $\mathcal{M}_1$ which is assigned to a link $j \in [m] \setminus [j_2 - 1]$ in the optimum strategy profile $\mathbf{L}$. This player has traffic at least w_{i_2}. Thus, for all $j \in \mathcal{M}_2$,

$$\frac{w_{i_2}}{c_j} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}). \tag{4.10}$$

So, for all $j \in \mathcal{M}_2$,

$$\begin{aligned}(k-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) &\overset{(4.9)}{\leq} \frac{\delta_{j}^{-i_2}(\mathbf{P}) + w_{i_2}}{c_j} \\ &\overset{(4.10)}{\leq} \frac{\delta_{j}^{-i_2}(\mathbf{P})}{c_j} + \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}), \end{aligned}$$

or equivalently

$$\delta_{j}^{-i_2}(\mathbf{P}) \geq (k-2) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}).$$

Therefore, for all $j \in \mathcal{M}_2$,

$$\begin{aligned}\delta_j(\mathbf{P}) &= \delta_{j}^{-i_2}(\mathbf{P}) + p_{i_2 j} \cdot w_{i_2} \\ &\geq \delta_{j}^{-i_2}(\mathbf{P}) \\ &\geq (k-2) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}), \end{aligned} \tag{4.11}$$

proving (b). Summing up all expected traffic $\delta_j(\mathbf{P})$ on links in $\mathcal{M}_2$, the total expected traffic of links in $\mathcal{M}_2$ is

$$\begin{aligned}\sum_{j \in \mathcal{M}_2} \delta_j(\mathbf{P}) &\overset{(4.11)}{\geq} (k-2) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_2} c_j \\ &\overset{(4.8)}{\geq} (k-2) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \cdot \rho \cdot C \cdot (k-2) \\ &= \rho \cdot C \cdot (k-2)^2 \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}). \end{aligned}$$

proving (c). In the optimum strategy profile $\mathbf{L}$ at most expected traffic

$$\sum_{j\in[j_2]} \delta_j(\mathbf{L}) \le \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \cdot \sum_{j\in[j_2]} c_j \tag{4.12}$$

can be assigned to links in $\mathcal{M}_1 \cup \mathcal{M}_2 = [j_2]$. So the remaining expected traffic on links in $\mathcal{M}_1 \cup \mathcal{M}_2$ which has to be assigned to other links in the optimal Nash equilibrium is at least

$$\begin{aligned}
&\sum_{j\in[j_2]} \delta_j(\mathbf{P}) - \sum_{j\in[j_2]} \delta_j(\mathbf{L}) \\
&\overset{(4.12)}{\ge} \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) + \sum_{j\in\mathcal{M}_2} \delta_j(\mathbf{P}) - \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \cdot \sum_{j\in[j_2]} c_j \\
&\overset{(4.4)(4.11)}{\ge} \left((k-1)\cdot \sum_{j\in\mathcal{M}_1} c_j + (k-2)\cdot \sum_{j\in\mathcal{M}_2} c_j - \sum_{j\in[j_2]} c_j \right) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&= \left((k-2)\cdot \sum_{j\in\mathcal{M}_1} c_j + (k-3)\cdot \sum_{j\in\mathcal{M}_2} c_j \right) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&\overset{(4.1)(4.8)}{\ge} ((k-2)\cdot\rho\cdot C + (k-3)\cdot(k-2)\cdot\rho\cdot C)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&= \rho\cdot C\cdot(k-2)^2 \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),
\end{aligned}$$

proving (d). This completes the proof of the claim for $l = 2$.
Now, assume inductively that for some integer $l \ge 3$ the claim holds for all integers not exceeding $(l-1)$. We will prove the claim for l.
Let w_{i_l} be the smallest traffic of a player i_l who assigns its traffic to a link in $\mathcal{M}_1 \cup \cdots \cup \mathcal{M}_{l-1}$ with positive probability, and let $s_l \in \mathcal{M}_1 \cup \cdots \cup \mathcal{M}_{l-1}$ be a link with $p_{i_l s_l} > 0$. By induction hypothesis we have

$$\begin{aligned}
&\sum_{j\in[j_{l-1}]} \delta_j(\mathbf{P}) - \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{L}) \\
&\qquad \ge \rho\cdot C\cdot(k-2)\cdot \prod_{j=2}^{l-1}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}).
\end{aligned} \tag{4.13}$$

This implies that there exists a set of links $\mathcal{M}_l = [j_l] \setminus [j_{l-1}]$, j_l minimal, with total capacity at least

$$\begin{aligned}
\sum_{j\in\mathcal{M}_l} c_j &\ge \frac{\sum_{j\in[j_{l-1}]} \delta_j(\mathbf{P}) - \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})} \\
&\overset{(4.13)}{\ge} \rho\cdot C\cdot(k-2)\cdot \prod_{j=2}^{l-1}(k-j),
\end{aligned} \tag{4.14}$$

proving (a). Moreover, since $\mathbf{P}$ is a Nash equilibrium, for all links $j \in \mathcal{M}_l$,

$$
\begin{aligned}
(k-(l-1)) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) &\leq \frac{\delta_{s_l}(\mathbf{P})}{c_{s_l}} \\
&\leq \frac{\delta_{s_l}^{-i_l}(\mathbf{P}) + w_{i_l}}{c_{s_l}} \\
&\leq \frac{\delta_j^{-i_l}(\mathbf{P}) + w_{i_l}}{c_j}. \qquad (4.15)
\end{aligned}
$$

By construction of $\mathcal{M}_l$, there exists a player assigned in $\mathbf{P}$ with positive probability to a link in $[j_{l-1}]$ which is assigned to a link $j \in [m] \setminus [j_l - 1]$ in the optimum strategy profile $\mathbf{L}$. This player has traffic at least w_{i_l}. Thus, for all $j \in \mathcal{M}_l$,

$$
\frac{w_{i_l}}{c_j} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \qquad (4.16)
$$

So, for all $j \in \mathcal{M}_l$,

$$
\begin{aligned}
(k-(l-1)) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) &\overset{(4.15)}{\leq} \frac{\delta_j^{-i_l}(\mathbf{P}) + w_{i_l}}{c_j} \\
&\overset{(4.16)}{\leq} \frac{\delta_j^{-i_l}(\mathbf{P})}{c_j} + \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),
\end{aligned}
$$

or equivalently

$$
\delta_j^{-i_l}(\mathbf{P}) \geq (k-l) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}).
$$

Therefore, for all $j \in \mathcal{M}_l$,

$$
\begin{aligned}
\delta_j(\mathbf{P}) &= \delta_j^{-i_l}(\mathbf{P}) + p_{i_l j} \cdot w_{i_l} \\
&\geq \delta_j^{-i_l}(\mathbf{P}) \\
&\geq (k-l) \cdot c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}), \qquad (4.17)
\end{aligned}
$$

proving (b). Summing up all expected traffic $\delta_j(\mathbf{P})$ on links in $\mathcal{M}_l$, the total expected traffic of links in $\mathcal{M}_l$ is

$$
\begin{aligned}
\sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{P}) &\overset{(4.17)}{\geq} (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j \\
&\overset{(4.14)}{\geq} \rho \cdot C \cdot (k-2) \cdot \prod_{j=2}^{l} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),
\end{aligned}
$$

proving (c). In the optimum strategy profile $\mathbf{L}$ at most traffic

$$
\sum_{j \in [j_l]} \delta_j(\mathbf{L}) \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \cdot \sum_{j \in [j_l]} c_j \qquad (4.18)
$$

can be assigned to links in $\mathcal{M}_1 \cup \cdots \cup \mathcal{M}_l = [j_l]$. So the remaining expected traffic on links in $\mathcal{M}_1 \cup \cdots \cup \mathcal{M}_l$ which has to be assigned to other links in the optimal solution is at least

$$\begin{aligned}
&\sum_{j\in[j_l]} \delta_j(\mathbf{P}) - \sum_{j\in[j_l]} \delta_j(\mathbf{L}) \\
&= \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{P}) - \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{L}) + \sum_{j\in\mathcal{M}_l} \delta_j(\mathbf{P}) - \sum_{j\in\mathcal{M}_l} \delta_j(\mathbf{L}) \\
&\overset{(4.17)(4.18)}{\geq} \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{P}) - \sum_{j\in[j_{l-1}]} \delta_j(\mathbf{L}) \\
&\quad +(k-l)\cdot \sum_{j\in\mathcal{M}_l} c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) - \sum_{j\in\mathcal{M}_l} c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&\overset{(4.13)}{\geq} \rho\cdot C\cdot (k-2)\cdot \prod_{j=2}^{l-1}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&\quad +(k-l-1)\cdot \sum_{j\in\mathcal{M}_l} c_j \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&\overset{(4.14)}{\geq} \rho\cdot C\cdot (k-2)\cdot \prod_{j=2}^{l-1}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&\quad +(k-l-1)\cdot \rho\cdot C\cdot (k-2)\cdot \prod_{j=2}^{l-1}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&= (k-l)\cdot \rho\cdot C\cdot (k-2)\cdot \prod_{j=2}^{l-1}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\
&= \rho\cdot C\cdot (k-2)\cdot \prod_{j=2}^{l}(k-j)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}),
\end{aligned}$$

proving (d). This completes the proof of the inductive claim. ■

(3.) Summing up the lower bounds on the expected loads over all links we get the lower bound $\sum_{j\in[m]} \delta_j(\mathbf{P}) < W$ on the total traffic W that is necessary for a Nash equilibrium $\mathbf{P}$ with $\mathsf{IC}(\mathbf{w},\mathbf{c},\mathbf{P}) = k$. Note that the strict inequality follows from the fact that we have at least one player with expected latency k. Using this lower bound, we now prove the upper bounds for the three cases of the theorem by showing that a larger upper bound implies $\frac{W}{C} > \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$, a contradiction. Note that $\Gamma^{-1}(\frac{1}{\rho})$ is also an upper bound on the ratio for $\rho \geq \frac{1}{37}$. However, in the ranges $\frac{1}{3} \leq \rho \leq 1$ and $\frac{1}{37} \leq \rho < \frac{1}{3}$ the given upper bounds are better. Now consider the three cases of the theorem:

(I) $\underline{\frac{1}{3} \leq \rho \leq 1}$: Assume $k \geq \frac{3}{2} + \sqrt{\frac{1}{\rho} - \frac{3}{4}}$. This implies $k \geq 2$ in the given range of ρ. Then

$$\begin{aligned} W &> \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) + \sum_{j\in\mathcal{M}_2} \delta_j(\mathbf{P}) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\big((k-1)+(k-2)^2\big) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\big(k^2-3\cdot k+3\big) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\ &\qquad\cdot\left(\left(\frac{3}{2}+\sqrt{\frac{1}{\rho}-\frac{3}{4}}\right)^2 - 3\cdot\left(\frac{3}{2}+\sqrt{\frac{1}{\rho}-\frac{3}{4}}\right)+3\right) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) \\ &\qquad\cdot\left(\frac{9}{4}+3\cdot\sqrt{\frac{1}{\rho}-\frac{3}{4}}+\frac{1}{\rho}-\frac{3}{4}-\frac{9}{2}-3\cdot\sqrt{\frac{1}{\rho}-\frac{3}{4}}+3\right) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\left(\frac{1}{\rho}\right) \\ &= C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \end{aligned}$$

(II) $\underline{\frac{1}{37}\leq\rho<\frac{1}{3}}$: Assume $k\geq 2+\sqrt[3]{\frac{1}{\rho}-2}$. This implies $k>3$ in the given range of ρ. Then,

$$\begin{aligned} W &> \sum_{j\in\mathcal{M}_1} \delta_j(\mathbf{P}) + \sum_{j\in\mathcal{M}_2} \delta_j(\mathbf{P}) + \sum_{j\in\mathcal{M}_3} \delta_j(\mathbf{P}) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\big((k-1)+(k-2)^2+(k-2)^2(k-3)\big) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\big(k-1+(k-2)^3\big) \\ &\overset{k>3}{>} \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\big(2+(k-2)^3\big) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot\left(2+\frac{1}{\rho}-2\right) \\ &= C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}). \end{aligned}$$

(III) $\underline{\rho<\frac{1}{37}}$: Assume $k\geq\Gamma^{-1}(\frac{1}{\rho})$. Using the facts that $\Gamma(x+1)=x\cdot\Gamma(x)$ for all $x\in\mathbb{R}$ and $\Gamma(x)\leq x$ for all $1\leq x\leq 3$, we get

$$\begin{aligned} W &> \sum_{j\in\mathcal{M}_{\lfloor k\rfloor-2}} \delta_j(\mathbf{P}) + \sum_{j\in\mathcal{M}_{\lfloor k\rfloor-1}} \delta_j(\mathbf{P}) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot(k-2)\cdot\left(\prod_{j=2}^{\lfloor k\rfloor-2}(k-j)+\prod_{j=2}^{\lfloor k\rfloor-1}(k-j)\right) \\ &> \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot(k-2)\cdot\left(\prod_{j=3}^{\lfloor k\rfloor-1}(k-j)+\prod_{j=2}^{\lfloor k\rfloor-1}(k-j)\right) \\ &\geq \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot(k-2)\cdot(\Gamma(k-2)+\Gamma(k-1)) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot(k-2)\cdot(\Gamma(k-2)+(k-2)\cdot\Gamma(k-2)) \\ &= \rho\cdot C\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\cdot(k-2)\cdot((k-1)\cdot\Gamma(k-2)) \end{aligned}$$

$$= \rho \cdot C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \cdot \Gamma(k)$$
$$\geq \rho \cdot C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \cdot \Gamma\left(\Gamma^{-1}\left(\frac{1}{\rho}\right)\right)$$
$$= C \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}).$$

In each of the cases we have $\frac{W}{C} > \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$. This is a contradiction to the fact that $\frac{W}{C} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$.

This completes the proof of the theorem. ■

Since $\frac{w_1}{c_1} \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$, we have $\rho \geq \frac{c_1}{C} \geq \frac{1}{m}$. Furthermore, $\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{P}) \geq \Lambda(\mathbf{P})$ holds for every strategy profile $\mathbf{P}$. Thus, from Theorem 4.10 we get the following corollaries:

Corollary 4.12. *Consider the model of arbitrary players and related links. Then for any instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and associated Nash equilibrium $\mathbf{P}$, the maximum expected latency $\Lambda(\mathbf{P})$ is bounded from above by*

$$\Lambda(\mathbf{P}) \leq \Gamma^{-1}\left(\frac{1}{\rho}\right) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \leq \Gamma^{-1}(m) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}).$$

Corollary 4.13. *Consider the model of arbitrary players and related links. Then for any instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and associated Nash equilibrium $\mathbf{P}$,*

$$\frac{\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{P})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})} \leq \Gamma^{-1}(m).$$

Corollary 4.14. *Consider the model of arbitrary players and related links. Then for any instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and associated Nash equilibrium $\mathbf{P}$,*

$$\mathsf{pPoA}_{\mathsf{MSP}} \leq \Gamma^{-1}(m).$$

We now introduce a pure strategy profile in Example 4.3.1 for which we show in Lemma 4.15 that it is a pure Nash equilibrium with certain properties. The pure Nash equilibrium will be used in Theorem 4.16 and Theorem 4.17 to prove that the upper bounds of $\Gamma^{-1}(\frac{1}{\rho})$ and $\Gamma^{-1}(m)$ are tight.

Example 4.3.1 *Let $k \in \mathbb{N}$, and consider the following instance $\langle \mathbf{w}, \mathbf{c} \rangle$ with associated pure strategy profile $\mathbf{L}$.*

- *There are k different classes of players:*
  - *Class U_1:* $|U_1| = k$ *players with traffic* 2^{k-1}
  - *Class U_i:* $|U_i| = 2^{i-1} \cdot (k-1) \prod_{j=1,\ldots,i-1}(k-j)$ *players with traffic* 2^{k-i} *for all* $2 \leq i \leq k$.
- *There are $k+1$ different classes of links:*
  - *Class P_0:* *One link with capacity* 2^{k-1}.
  - *Class P_1:* $|P_1| = |U_1| - 1$ *links with capacity* 2^{k-1}.
  - *Class P_i:* $|P_i| = |U_i|$ *links with capacity* 2^{k-i} *for all* $2 \leq i \leq k$.

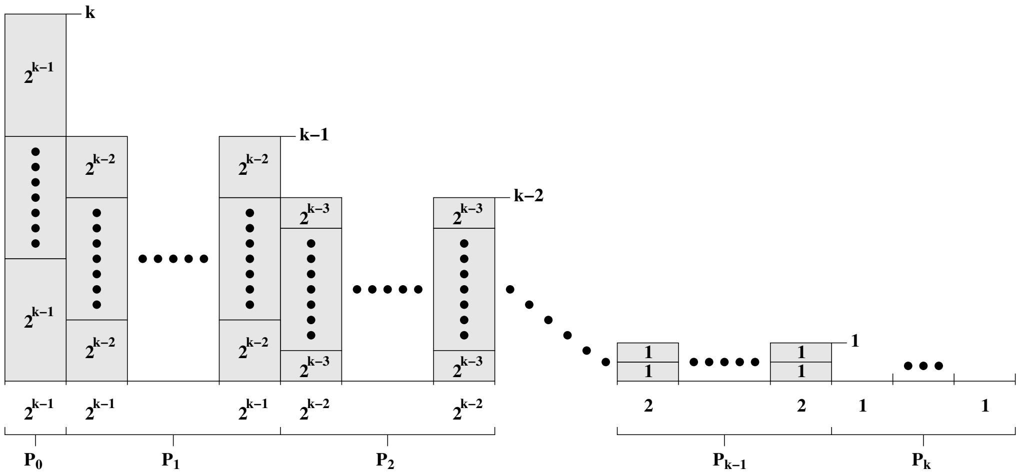

Fig. 4.2. Instance and associated strategy profile **L** from Example 4.3.1

- *Consider the following strategy profile* $\mathbf{L}$*:*
  - *Class* P_0*: All players in* U_1 *are assigned to this link.*
  - *Class* P_i*: On each link in* P_i *there are* $2(k-i)$ *players from* U_{i+1}*, respectively, for all* $1 \leq i \leq k-1$*.*
  - *Class* P_k*: The links from* P_k *remain empty.*

Lemma 4.15. *Consider the instance* $\langle \mathbf{w}, \mathbf{c} \rangle$ *and associated pure strategy profile* $\mathbf{L}$ *given in Example 4.3.1. Then*

(a) $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = 1$*, and*
(b) the strategy profile $\mathbf{L}$ *is a pure Nash equilibrium with* $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = k$*.*

Proof. We will prove (a) and (b) separately.

(a) In an optimum strategy profile the player with largest traffic must be assigned to some link. Thus,
$$\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \geq \frac{w_1}{c_1} = 1.$$
Now, consider the pure strategy profile $\mathbf{L}'$ that assigns one player from U_1 to each link in $P_0 \cup P_1$ and one player from U_i to each link in P_i for all $2 \leq i \leq k$. Then for all links $j \in [m]$, we have $\frac{\delta_j(\mathbf{L}')}{c_j} = 1$, so that
$$\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}') = \max_{j \in [m]} \frac{\delta_j(\mathbf{L}')}{c_j} = 1.$$
It follows that $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = 1$.

(b) Clearly, $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = k$. We now show, that $\mathbf{L}$ is a Nash equilibrium. In the pure strategy profile $\mathbf{L}$ the latency of each link $j \in P_i$ where $i \in [k]_0$ is
$$\Lambda_j(\mathbf{L}) = \frac{\delta_j(\mathbf{L})}{c_j} = k - i. \tag{4.19}$$

By way of contradiction, assume that the pure strategy profile $\mathbf{L}$ is not a Nash equilibrium. Then, there exists some player s with traffic w_s that is not satisfied in $\mathbf{L}$. Assume that player s can improve by moving from some link $j_1 \in P_{i_1}$ to some other link $j_2 \in P_{i_2}$. Thus,

$$\frac{\delta_{j_1}(\mathbf{L})}{c_{j_1}} > \frac{\delta_{j_2}(\mathbf{L}) + w_s}{c_{j_2}}. \tag{4.20}$$

By (4.19) it must hold that $i_1 < i_2$. We proceed by case analysis:
First assume $i_1 = 0$: Then $w_s = 2^{k-1}$ and

$$\begin{aligned} k &\stackrel{(4.19)}{=} \frac{\delta_{j_1}(\mathbf{L})}{c_{j_1}} \\ &\stackrel{(4.20)}{>} \frac{\delta_{j_2}(\mathbf{L}) + w_s}{c_{j_2}} \\ &\stackrel{(4.19)}{=} (k - i_2) + \frac{2^{k-1}}{2^{k-i_2}} \\ &= (k - i_2) + 2^{i_2 - 1} \\ &\stackrel{i_2 \geq 1}{\geq} (k - i_2) + i_2 \\ &= k, \end{aligned}$$

a contradiction.
Now assume $i_1 \geq 1$: Then $w_s = 2^{k-i_1-1}$ and

$$\begin{aligned} k - i_1 &\stackrel{(4.19)}{=} \frac{\delta_{j_1}(\mathbf{L})}{c_{j_1}} \\ &\stackrel{(4.20)}{>} \frac{\delta_{j_2}(\mathbf{L}) + w_s}{c_{j_2}} \\ &\stackrel{(4.19)}{=} (k - i_2) + \frac{2^{k-i_1-1}}{2^{k-i_2}} \\ &= (k - i_2) + 2^{i_2 - i_1 - 1} \\ &\stackrel{i_2 - i_1 \geq 1}{\geq} (k - i_2) + i_2 - i_1 \\ &= k - i_1, \end{aligned}$$

a contradiction. It follows that $\mathbf{L}$ is a pure Nash equilibrium.

This completes the proof of the lemma. ∎

We are now ready to prove two lower bounds on the individual price of anarchy. Our lower bounds hold for the case of pure Nash equilibria.

Theorem 4.16. *For each $k \in \mathbb{N}$ there exists an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and an associated pure Nash equilibrium $\mathbf{L}$ with*

$$k = \frac{\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})} \geq \Gamma^{-1}\left(\frac{1}{\rho}\right) - 1.$$

Proof. Consider the instance $\langle \mathbf{w}, \mathbf{c} \rangle$ from Example 4.3.1 with associated pure strategy profile $\mathbf{L}$. By Lemma 4.15, $\mathbf{L}$ is a pure Nash equilibrium for $\langle \mathbf{w}, \mathbf{c} \rangle$, $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = 1$ and $\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = k$. We now prove that $\Gamma^{-1}(\frac{1}{\rho}) - 1$ is a lower bound on k. We have

$$\rho = \frac{|P_0 \cup P_1| \cdot 2^{k-1}}{|P_0 \cup P_1| \cdot 2^{k-1} + \sum_{i \in [2,k]} |P_i| \cdot 2^{k-i}}.$$

This implies

$$\begin{aligned}
\frac{1}{\rho} &= \frac{1}{|P_0 \cup P_1| \cdot 2^{k-1}} \cdot \left(|P_0 \cup P_1| \cdot 2^{k-1} + \sum_{i \in [2,k]} |P_i| \cdot 2^{k-i} \right) \\
&= \frac{1}{k \cdot 2^{k-1}} \cdot \left(k \cdot 2^{k-1} + \sum_{i \in [2,k]} \left(2^{i-1} \cdot (k-1) \prod_{j \in [i-1]} (k-j) \right) \cdot 2^{k-i} \right) \\
&= \frac{1}{k \cdot 2^{k-1}} \cdot \left(k \cdot 2^{k-1} + 2^{k-1} \cdot (k-1) \sum_{i \in [2,k]=2} \prod_{j \in [i-1]} (k-j) \right) \\
&< 1 + \sum_{i \in [2,k]} \prod_{j \in [i-1]} (k-j) \\
&< 1 + (k-1) \cdot (k-1)! \\
&= 1 + k! - (k-1)! \\
&\overset{k \geq 1}{\leq} k! \\
&= \Gamma(k+1).
\end{aligned}$$

This yields $k \geq \Gamma^{-1}\left(\frac{1}{\rho}\right) - 1$, which completes the proof of the theorem. ■

Theorem 4.17. *For each $k \in \mathbb{N}$ there exists an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ and an associated pure Nash equilibrium $\mathbf{L}$ with*

$$k = \frac{\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})} \geq \Gamma^{-1}(m) \cdot (1 + o(1)).$$

Proof. Consider the instance $\langle \mathbf{w}, \mathbf{c} \rangle$ from Example 4.3.1 with associated pure strategy profile $\mathbf{L}$. By Lemma 4.15, $\mathbf{L}$ is a pure Nash equilibrium for $\langle \mathbf{w}, \mathbf{c} \rangle$, $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = 1$ and $\mathsf{IC}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) = k$. Moreover, we have

$$\begin{aligned}
m = \sum_{i \in [k]_0} |P_i| &= k + (k-1) \cdot \sum_{i \in [2,k]} 2^{i-1} \cdot \prod_{j \in [i-1]} (k-j) \\
&= k + (k-1) \cdot 2^{k-1} \cdot (k-1)! \cdot \left(1 + \sum_{i \in [2,k-1]} \frac{1}{2^{k-i}} \cdot \frac{1}{(k-i)!} \right) \\
&\leq k + (k-1) \cdot 2^k \cdot (k-1)! \\
&\leq 2^k \cdot k! \\
&\leq \alpha \cdot k^k,
\end{aligned}$$

for a constant $\alpha \in \mathbb{R}^+$. We define

$$r = \alpha \cdot k^k.$$

Since

$$\begin{aligned}\log(r) &= k \cdot \log(k) + \log(\alpha) \\ &= k \cdot \log(k) \cdot (1 + o(1))\end{aligned}$$

and

$$\begin{aligned}\log(\log(r)) &= \log(k) + \log(\log(k)) + o(1) \\ &= \log(k) \cdot (1 + o(1)),\end{aligned}$$

this implies

$$\begin{aligned}\Gamma^{-1}(m) &\leq \Gamma^{-1}(r) \\ &= \frac{\log(r)}{\log(\log(r))} \cdot (1 + o(1)) \\ &= k \cdot (1 + o(1)).\end{aligned}$$

This completes the proof of the claim. ■

4.4 Restricted Strategy Sets

For all our results so far we have assumed that the strategy sets of the players are unrestricted. We will now drop this assumption and consider routing games on parallel links with *restricted strategy sets*. Here, each player $i \in [n]$ is only allowed to assigned its traffic to links in its strategy set S_i, where $S_i \subseteq [m]$. In Section 4.4.1 we study the problem of computing a pure Nash equilibrium, while Section 4.4.2 deals with the pure price of anarchy.

4.4.1 Computation of Pure Nash Equilibria

For the model arbitrary players with restricted strategy sets and related links it is not known whether a pure Nash equilibrium can be computed in polynomial time. However, for the case of *identical* players, a result of Milchtaich [77] implies the existence of a polynomial-time algorithm to perform this task.

Furthermore, for the model of arbitrary players with restricted strategy sets and identical links, Gairing et al. [41] presented a polynomial-time algorithm, called NASHIFY-RESTRICTED, to compute a pure Nash equilibrium.

In the following, we will present NASHIFY-RESTRICTED, which combines ideas from blocking flows and the generic PREFLOW-PUSH algorithm [3]. Here, we assume that all players' traffic is integer. In Section 4.4.1.1, we will first show, how to represent a pure strategy profile by a residual network. Afterwards, in Section 4.4.1.2, we introduce a blocking flow algorithm, called UNSPLITTABLE-BLOCKING-FLOW. In Section 4.4.1.3, we show how UNSPLITTABLE-BLOCKING-FLOW can be used to convert a given strategy profile into a pure Nash equilibrium with non-increased makespan social cost.

4.4.1.1 Residual Network Representation

We introduce a residual network $G_{\mathbf{L}}$ representing a pure strategy profile $\mathbf{L}$.

Definition 4.18. *Given a pure strategy profile* $\mathbf{L} = (\ell_1, \ldots, \ell_n)$*, we define a directed bipartite graph* $G_{\mathbf{L}} = (V, E_{\mathbf{L}})$*, where* $V = M \cup U$ *such that each link is represented by a node in* M *and each player is represented by a node in* U*. Furthermore,* $E_{\mathbf{L}} = E^1_{\mathbf{L}} \cup E^2_{\mathbf{L}}$ *with*

$$E^1_{\mathbf{L}} = \{(u, v) : u \in M, v \in U, u = \ell_v\}, \text{ and}$$
$$E^2_{\mathbf{L}} = \{(u, v) : u \in U, v \in M, v \in S_u \setminus \{\ell_u\}\}.$$

For an arbitrary integer w, we use the graph $G_{\mathbf{L}}$ from Definition 4.18 to define a graph $G_{\mathbf{L}}(w)$ where V remains the same, but from $E_{\mathbf{L}}$ we now only consider edges $E_{\mathbf{L}}(w) = E_{\mathbf{L}} \setminus \{(u, v) : u \in U, v \in V, w_u > w\}$. This implies that players u with $w_u > w$ remain assigned to their links. We use $G_{\mathbf{L}}$ for $G_{\mathbf{L}}(w)$ whenever w is clear from context.

4.4.1.2 Unsplittable Blocking Flow

We now introduce a blocking flow algorithm, called Unsplittable-Blocking-Flow, which will be extensively used by our Nashification algorithm in Section 4.4.1.3. Unsplittable-Blocking-Flow combines ideas from blocking flows with the idea of pushing players without splitting them.

To control our blocking flow algorithm we use two integer parameters a and w. Here, w will be used to refer to a certain traffic size, and a will be determined by binary search. For every integer a and traffic size w we partition the set of links M into three subsets:

$$M^- = \{j \in M \mid \delta_j(\mathbf{L}) \leq a\}$$
$$M^0 = \{j \in M \mid a + 1 \leq \delta_j(\mathbf{L}) \leq a + w\}$$
$$M^+ = \{j \in M \mid \delta_j(\mathbf{L}) \geq a + w + 1\}$$

In this setting, we do not have a dedicated source or sink. However, at each time nodes in M^+ and M^- can be interpreted as source and sink nodes, respectively. Note, that those sets change over time.

Roughly speaking, algorithm Unsplittable-Blocking-Flow shifts players so that the latencies of links from M^- are never decreased, the latencies of links from M^+ are never increased, and links from M^0 remain in M^0. Our algorithm is controlled by a height function $h : V \to \mathbb{N}_0$ with $h(j) = \text{dist}_{G_{\mathbf{L}}}(j, M^-)$ for all $j \in V$. We call an edge (u, v) *admissible*, if $h(u) = h(v) + 1$. In an *admissible path*, all edges are admissible. For each node $j \in V$ with $0 < h(j) < \infty$, define `Suc(j)` to be the set of successors of node j; this is the set of nodes to which j has an admissible edge, so that

$$\texttt{Suc(j)} = \{i \in V : (j, i) \in E_{\mathbf{L}} \text{ and } h(j) = h(i) + 1\}.$$

Note that Suc(j) also defines the set of admissible edges leaving j. Let suc(j) be the first node in a list implementation of the set Suc(j). We proceed to define:

Definition 4.19. *A link $j \in M$ with $0 < h(j) < \infty$ is called helpful if $\delta_j(\mathbf{L}) \geq a + 1 + w_{\mathtt{suc(j)}}$.*

Lemma 4.20 (Gairing et al. [41]). *Let v_0 be a helpful link of minimum height. Then there exists a sequence $v_0, \ldots, v_r$, where $v_{2i} \in M$ for all $0 \leq i \leq r/2$ and $v_{2i+1} \in U$ for all $0 \leq i < r/2$ such that:*

(1) $(v_i, v_{i+1}) \in E_{\mathbf{L}}$ *and* $h(v_i) = h(v_{i+1}) + 1$.
(2) $\delta_{v_0}(\mathbf{L}) \geq a + 1 + w_{\mathtt{suc(v_0)}}$.
(3) $a + 1 \leq \delta_{v_{2i}}(\mathbf{L}) + w_{\mathtt{suc(v_{2i-2})}} - w_{\mathtt{suc(v_{2i})}} \leq a + w, \forall 0 < i < r/2$.
(4) $\delta_{v_r}(\mathbf{L}) + w_{\mathtt{suc(v_{r-2})}} \leq a + w$.

We are now ready to present the algorithm Unsplittable-Blocking-Flow. The algorithm is depicted in Figure 4.3. Initially, the height function h is computed as the distance in $G_{\mathbf{L}}$ of each node to the set M^- of nodes. Then, the algorithm proceeds in phases. In each phase first the minimum height $d = h(v)$ of a node $v \in M^+$ is computed. Inside each phase, we do not update the height function, but we successively choose a helpful link v of minimum height and we push players along the helpful path induced by v and adjust the pure strategy profile accordingly. In order to update $G_{\mathbf{L}}$ we have to change the direction of two arcs for each player push. The phase ends when there exists no further admissible path from an node $v \in M^+$ with $h(v) = d$ to some node in M^-. Before the new phase starts, we recompute h, and we check whether we need to start a new phase or not. Unsplittable-Blocking-Flow stops when either $M^- = \emptyset$ or for all $v \in M^+$ we have $h(v) = \infty$.

Unsplittable-Blocking-Flow($\mathbf{L}, a, w$)
Input: pure strategy profile $\mathbf{L}$
  positive integers a, w
Output: pure strategy profile $\mathbf{L}'$
compute h;
while $M^- \neq \emptyset$ **and** $\exists v \in M^+ : h(v) < \infty$ **do**
  $d \leftarrow \min_{v \in M^+}(h(v))$;
  while $\exists$ admissible path from $v \in M^+, h(v) = d$ to M^- **do**
    choose helpful link v of minimum height;
    push players along helpful path defined by v;
    update $\mathbf{L}, G_{\mathbf{L}}$;
  end while
  recompute h;
end while
return $\mathbf{L}$;

Fig. 4.3. Unsplittable-Blocking-Flow

Gairing et al. [41] showed that a call to Unsplittable-Blocking-Flow does not increase the load on any link in M^+, does not decrease the load on any link in M^-, and that links in M^0 remain in M^0 (Lemma 4.21). This implies that UNSPLITTABLE-BLOCKING-FLOW does not increase the maximum load and not decreases the minimum load on a link (Corollary 4.22). Furthermore, they showed properties on the pure strategy profile computed by UNSPLITTABLE-BLOCKING-FLOW (Lemma 4.23) and the running time of UNSPLITTABLE-BLOCKING-FLOW (Theorem 4.24).

Lemma 4.21 (Gairing et al. [41]). *Let $\mathbf{L}'$ be the pure strategy profile computed by* UNSPLITTABLE-BLOCKING-FLOW*($\mathbf{L}, a, w$). Then,*

(1) $\delta_j(\mathbf{L}') \geq \delta_j(\mathbf{L})$ for each link $j \in M^-(\mathbf{L})$.
(2) $a + 1 \leq \delta_j(\mathbf{L}') \leq a + w$ for each link $j \in M^0(\mathbf{L})$.
(3) $\delta_j(\mathbf{L}') \leq \delta_j(\mathbf{L})$ for each node $j \in M^+(\mathbf{L})$.

Corollary 4.22 (Gairing et al. [41]). *Let $\mathbf{L}'$ be the pure strategy profile computed by* UNSPLITTABLE-BLOCKING-FLOW*($\mathbf{L}, a, w$). Then,*

$$\max_{j\in[m]} \delta_j(\mathbf{L}') \leq \max_{j\in[m]} \delta_j(\mathbf{L}), \text{ and}$$
$$\min_{j\in[m]} \delta_j(\mathbf{L}') \geq \min_{j\in[m]} \delta_j(\mathbf{L}).$$

Lemma 4.23 (Gairing et al. [41]). *For the pure strategy profile $\mathbf{L}' = (\ell'_1, \ldots, \ell'_n)$ computed by* UNSPLITTABLE-BLOCKING-FLOW*($\mathbf{L}, a, w$) one of the following conditions holds:*

(1) $M^-(\mathbf{L}') = \emptyset$.
(2) $M^+(\mathbf{L}') = \emptyset$.
(3) There exists some set of links $B \subset [m]$ such that
  a) $\delta_j(\mathbf{L}') \geq a + 1$ for all $j \in B$, and
  b) $\delta_j(\mathbf{L}') \leq a + w$ for all $j \in [m] \setminus B$, and
  c) $\ell'_i \in B \Rightarrow S_i \subseteq B$ for all $i \in [n]$ with $w_i \leq w$.

Theorem 4.24 (Gairing et al. [41]). UNSPLITTABLE-BLOCKING-FLOW *can be implemented to run in $O(mA)$ time, where $A = \sum_{i\in[n]} |S_i|$.*

4.4.1.3 Nashification

We now describe how UNSPLITTABLE-BLOCKING-FLOW can be used to convert any pure strategy profile into a pure Nash equilibrium with non-increased social cost.

Our Nashification algorithm, called NASHIFY-RESTRICTED, first finds a pure strategy profile satisfying all players with traffic w_1 by recursively applying UNSPLITTABLE-BLOCKING-FLOW. In this recursive procedure, called RECURSIVEUBF, we make extensive use of Lemma 4.23.

We then fix the pure strategy profile of all players with traffic w_1 and proceed with the next smaller traffic while making sure that all fixed players stay satisfied. To make sure that all fixed players stay satisfied, we introduce lower and upper bounds on the load of the links, such that the load of each link is always in its bounds, the lower bound only increases and the upper bound only decreases. This is done until all players are satisfied. In order to achieve this, NASHIFY-RESTRICTED makes extensive use of algorithm UNSPLITTABLE-BLOCKING-FLOW.

We now proceed with a detailed description of our Nashification algorithm. In the following, we denote $w = w_i$ for some player $i \in [n]$.

RecursiveUBF. We first turn our attention to RECURSIVEUBF, which is depicted in Figure 4.4. If $l \leq \delta_j(\mathbf{L}) \leq u + w$ for all links $j \in B$ prior to a call

RECURSIVEUBF$(B, \mathbf{L}(B), [l,u], w)$
Input: A set of links B, a pure strategy profile $\mathbf{L}(B)$, an interval $[l,u]$ and a traffic size w.
Output: A a pure strategy profile $\mathbf{L}'(B)$.
$a \leftarrow \lceil (l+u)/2 \rceil$;
if $a = u$ **then**
    return $\mathbf{L}(B)$
end if
$\mathbf{L}'(B) \leftarrow$ UNSPLITTABLE-BLOCKING-FLOW$(\mathbf{L}(B), a, w)$;
if $M^-(\mathbf{L}') = \emptyset$ **and** $M^+(\mathbf{L}') \neq \emptyset$ **then**
    $\mathbf{L}(B) \leftarrow$ RECURSIVEUBF$(B, \mathbf{L}(B), [a,u], w)$;
else if $M^-(\mathbf{L}') \neq \emptyset$ **and** $M^+(\mathbf{L}') = \emptyset$ **then**
    $\mathbf{L}(B) \leftarrow$ RECURSIVEUBF$(B, \mathbf{L}(B), [l,a], w)$;
else if $M^-(\mathbf{L}') \neq \emptyset$ **and** $M^+(\mathbf{L}') \neq \emptyset$ **then**
    split B (according to Lemma 4.23 (3)) into sets B' and $\overline{B'}$;
    $\mathbf{L}(B') \leftarrow$ RECURSIVEUBF$(B', \mathbf{L}(B'), [a,u], w)$;
    $\mathbf{L}(\overline{B'}) \leftarrow$ RECURSIVEUBF$(\overline{B'}, \mathbf{L}(\overline{B'}), [l,a], w)$;
    $\mathbf{L}(B) \leftarrow \mathbf{L}(B') \cup \mathbf{L}(\overline{B'})$;
end if
return $\mathbf{L}(B)$;

Fig. 4.4. RECURSIVEUBF

to RECURSIVEUBF$(B, \mathbf{L}(B), [l,u], w)$, then it computes a pure strategy profile, where no player with traffic w that is assigned to some link in B can improve by moving to some other link in B. By a series of calls to UNSPLITTABLE-BLOCKING-FLOW$(\mathbf{L}(B), a, w)$ we compute a pure strategy profile where M^- and M^+ are either both empty or both non-empty. Parameter a is chosen by binary search $a \in [l,u], a \in \mathbb{N}$, as follows: If UNSPLITTABLE-BLOCKING-FLOW returns a pure strategy profile with $M^- = \emptyset$ and $M^+ \neq \emptyset$, then we increase a. On the other hand, if UNSPLITTABLE-BLOCKING-FLOW returns a pure strategy profile with $M^- \neq \emptyset$ and $M^+ = \emptyset$, then we decrease a.

If after the binary search, $M^- = \emptyset$ and $M^+ = \emptyset$, then we have computed a pure strategy profile where all players with traffic at least w are satisfied. If neither $M^- = \emptyset$ nor $M^+ = \emptyset$ it follows that condition (3) from Lemma 4.23 holds. Define B' as the set of links still reachable from M^+ and let $\overline{B'}$ be the complement of B' in B. In this case we split our instance into two parts. One part with all links in B' and all players that are currently assigned to a link in

B', the other part holds the complement. Whenever B is split into B' and $\overline{B'}$, condition (3) from Lemma 4.23 implies that no player v with $w_v \leq w$, assigned to a link in B', has a link from $\overline{B'}$ in its strategy set.

We recursively proceed with the binary search on a in both parts of the instance. For the part that corresponds to B', we increase a, while in the other part we decrease a. The recursive splitting of B (line 11) defines a partition of the links into sets $B_1, \ldots, B_p$. At the end, all parts $B_1, \ldots, B_p$ are put together to form $\mathbf{L}(B)$.

For each $B_k, k \in [p]$, define a lower bound $\mathsf{Low}(B_k)$ on the load of all links from B_k as the last value for a after the binary search on a in B_k. This implies:

Lemma 4.25 (Gairing et al. [41]). *If $l \leq \delta_j(\mathbf{L}) \leq u + w$ for all $j \in B$ prior to a call to* RecursiveUBF*, then* RecursiveUBF$(B, \mathbf{L}(B), [l, u], w)$ *returns a pure strategy profile $\mathbf{L}'(B)$ of players in B, a partition of B into p sets $B_1, \ldots, B_p$ for some p, and (implicit) numbers $\mathsf{Low}(B_k)$ for $k \in [p]$, such that:*

(1) $u \geq \mathsf{Low}(B_1) > \ldots > \mathsf{Low}(B_p) \geq l$ for all $k \in [p]$.
(2) $\mathsf{Low}(B_k) \leq \delta_j \leq \mathsf{Low}(B_k) + w$ for all $j \in B_k$ and for all $k \in [p]$.
(3) No player u with $w_u \leq w$ assigned to a link in B_k has a link from B_ℓ in its strategy set, if $\ell > k$.

By the postconditions of Lemma 4.25 all players with traffic w are satisfied in the pure strategy profile computed by RecursiveUBF. In order to keep these players satisfied, we have to ensure that in further computations the lower bounds only increase and the upper bounds only decrease. We denote the upper bound by $\mathsf{Up}(B_k)$ for all links from B_k, and in coincidence with (2) we set $\mathsf{Up}(B_k) = \mathsf{Low}(B_k) + w$.

Nashify-Restricted. We are ready to present algorithm Nashify-Restricted that converts any given pure strategy profile $\mathbf{L}$ into a pure Nash equilibrium $\mathbf{L}'$ with non-increased social cost. Let $\widetilde{w}_1 > \ldots > \widetilde{w}_r$ be all different player traffic from $w_1, \ldots, w_n$. The idea is to compute a sequence of pure strategy profiles $\mathbf{L}_0, \ldots, \mathbf{L}_r$ such that $\mathbf{L}_0 = \mathbf{L}$, $\mathbf{L}_r = \mathbf{L}'$ and such that for all pure strategy profiles $\mathbf{L}_i$ with $1 \leq i \leq r$, all players u with $w_u \geq \widetilde{w}_i$ are satisfied. We call the computation of $\mathbf{L}_i$ from $\mathbf{L}_{i-1}$ *stage* i. The aim in stage i is to compute a pure strategy profile $\mathbf{L}_i$ from $\mathbf{L}_{i-1}$ such that in $\mathbf{L}_i$ all players u with $w_u \geq \widetilde{w}_i$ are satisfied.

Figure 4.5 shows the high-level structure of our Nashification algorithm. It first uses the procedure RecursiveUBF to compute a pure strategy profile $\mathbf{L}_1$, where all players with traffic $\widetilde{w}_1$ are satisfied. Afterwards we iteratively satisfy players with traffic $\widetilde{w}_2, \ldots, \widetilde{w}_r$ making sure that players with larger traffic remain satisfied (lines 2-6). We do this by executing Sweep over the sets of active links. In the following, we define what we mean by sets of active links, and we describe how a Sweep over these sets of active links is executed.

Lemma 4.25 implies that after stage 1, all players with traffic $\widetilde{w}_1$ are satisfied. Furthermore, the links are partitioned into p_1 sets $B_1, \ldots, B_{p_1}$ with $\mathsf{Up}(B_k) = \mathsf{Low}(B_k) + \widetilde{w}_1$ for all $k \in [p_1]$, and no player u with $w_u \leq \widetilde{w}_1$, that is assigned to a link from B_k can be assigned to a link from B_ℓ when $k < \ell$.

NASHIFY-RESTRICTED($\mathbf{L}_0$)

▷ stage 1:

1: $\mathbf{L}_1 \leftarrow$RECURSIVEUBF$([m], \mathbf{L}_0, [0, \max_j \delta_j(\mathbf{L}_0)], \widetilde{w}_1)$

▷ stages 2,. . . , r:

2: **for** $i \leftarrow 2$ to r **do**
3:     **while** there are sets of active links **do**
4:         execute SWEEP over the active links;
5:     **end while**

▷ $\mathbf{L}_i$ is the current pure strategy profile:

6: **end for**
7: **return** $\mathbf{L}_r$

Fig. 4.5. NASHIFY-RESTRICTED

We now describe stage $i > 1$ (lines 2-6 in Figure 4.5). The lower bound on the load of a link only increases and the upper bound only decreases. This implies that fixed players remain satisfied. At the beginning of stage i, we have a pure strategy profile $\mathbf{L}_{i-1}$, where the links are partitioned into p_{i-1} sets $B_1, \ldots, B_{p_{i-1}}$ with $\mathsf{Up}(B_k) = \mathsf{Low}(B_k) + \widetilde{w}_{i-1}$, for all $k \in [p_{i-1}]$, and no player u, that is assigned to a link from B_k can be assigned to a link from B_ℓ when $k < \ell$.

During each stage i, we always maintain a pure strategy profile where the links are partitioned into q sets $C_1, \ldots, C_q$ for some q. They are ordered such that $\mathsf{Up}(C_k) > \mathsf{Up}(C_{k+1})$ and $\mathsf{Low}(C_k) \geq \mathsf{Low}(C_{k+1})$ for all k with $1 \leq k < q$.

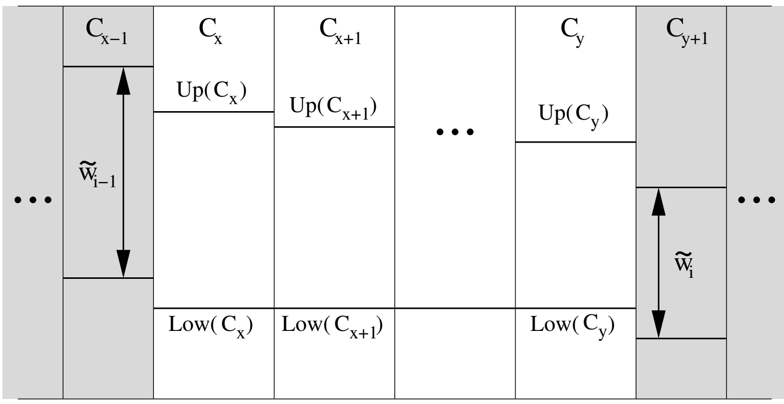

Fig. 4.6. Sets of active links in stage i at the beginning of a sweep

At the beginning of a SWEEP, we have three classes of sets (see Figure 4.6):

- Some sets of links $C_k, 1 \leq k < x$, have not been considered yet and fulfill $\mathsf{Up}(C_k) - \mathsf{Low}(C_k) = \widetilde{w}_{i-1}$.
- Moreover, some sets of links $C_k, q \geq k > y$, have been *done in stage* i already and fulfill $\mathsf{Up}(C_k) - \mathsf{Low}(C_k) = \widetilde{w}_i$.
- Finally, we have sets $C_x, \ldots, C_y$ of *active* links, with $\widetilde{w}_i < \mathsf{Up}(C_k) - \mathsf{Low}(C_k) \leq \widetilde{w}_{i-1}$ and $\mathsf{Low}(C_k) = \mathsf{Low}(C_y)$ for all $k \in [x, y]$.

Initially, $C_j = B_j$ for all $1 \leq j \leq p_{i-1}$, the links from $C_{p_{i-1}}$ are active, and the remaining links have not been considered. During a SWEEP, the number of partitions q may change. We will see in Lemma 4.28 that at the beginning of each SWEEP, the *sweep property* introduced below holds:

Definition 4.26 (Sweep Property during stage i).

(1) There is a partition of the links into q sets $C_1, \dots, C_q$ for some q with $\mathsf{Low}(C_1) \geq \dots \geq \mathsf{Low}(C_q)$ and $\mathsf{Up}(C_1) > \dots > \mathsf{Up}(C_q)$.
(2) If link $j \in C_k$, then $\mathsf{Low}(C_k) \leq \delta_j \leq \mathsf{Up}(C_k)$.
(3) No player u with $w_u \leq \widetilde{w}_i$ that is assigned to a link in C_k has a link from C_ℓ in its strategy set S_u, if $\ell > k$.
(4) There exist integers x, y with $1 \leq x \leq y \leq q$ and
a) $\mathsf{Up}(C_k) - \mathsf{Low}(C_k) = \widetilde{w}_{i-1}$ for $1 \leq k < x$,
b) $\mathsf{Up}(C_k) - \mathsf{Low}(C_k) = \widetilde{w}_i$ for $y < k \leq q$, and
c) $\widetilde{w}_i < \mathsf{Up}(C_k) - \mathsf{Low}(C_k) \leq \widetilde{w}_{i-1}$ and $\mathsf{Low}(C_k) = \mathsf{Low}(C_y)$ for all $x \leq k \leq y$.

We use the definition of sweep property to define active links.

Definition 4.27. *Let x, y be as in Definition 4.26. Then, a link $j \in C_k$, $x \leq k \leq y$, is called* active *and a link $j \in C_k$, $y < k \leq p$, is called* done *in stage i.*

A SWEEP is shown in Figure 4.7 and works on active links as follows: At the beginning of SWEEP, the sweep property holds. The aim of SWEEP is to process links in C_y such that they do not have to be considered again in this stage, or to make all links in C_{x-1} active by increasing the lower bound of all active links to $\mathsf{Low}(C_{x-1})$. In order to preserve the structure of our pure strategy profile, we choose $a = \min\{\mathsf{Up}(C_y) - \widetilde{w}_i, \mathsf{Low}(C_{x-1})\}$. We insert all sets into a list $\mathcal{L}$ such that $\mathcal{L} = [C_x, \dots, C_y]$. Then, as long as there are at least two sets in $\mathcal{L}$, we do the following: We extract the first element, say D_1, of $\mathcal{L}$ and apply UNSPLITTABLE-BLOCKING-FLOW to the sub-instance defined by the set D_1. UNSPLITTABLE-BLOCKING-FLOW$(\mathbf{L}(D_1), a, \widetilde{w}_i)$ returns a pure strategy profile $\mathbf{L}'$, where one of the following conditions holds:

1. $M^+(\mathbf{L}') = \emptyset$: In this case, all links in D_1 have load at most $a + \widetilde{w}_i$, and Corollary 4.22 implies that this property is preserved. Let D_2 be the next element in $\mathcal{L}$. Before the call, $\mathsf{Up}(D_1) > \mathsf{Up}(D_2) > a + \widetilde{w}_i$ was true. After the call, the loads of all links in D_1 are bounded by $a + \widetilde{w}_i$. So, by setting $\mathsf{Up}(D_1) \leftarrow \mathsf{Up}(D_2)$, we get a new upper bound on the loads of the links in D_1, and we fulfill the requirement that upper bounds can be only decreased. D_1 and D_2 are merged, and the union of both sets is inserted into $\mathcal{L}$. This way, the number of sets in the list is decreased by 1.
2. $M^-(\mathbf{L}') = \emptyset$ and $M^+(\mathbf{L}') \neq \emptyset$: In this case, all links in D_1 have load at least a, and Corollary 4.22 implies that this property is preserved. Thus, we are allowed to set $\mathsf{Low}(D_1) \leftarrow a$. We are done with D_1 during this execution of SWEEP.

Require: $\mathcal{L} = [C_x, \dots, C_y]$ is a list of the sets of *active* links
$a \leftarrow \min\{\mathsf{Up}(C_y) - \widetilde{w}_i, \mathsf{Low}(C_{x-1})\}$;
while $|\mathcal{L}| \geq 2$ **do**
  $D_1 \leftarrow \mathsf{ExtractFirst}(\mathcal{L})$;
  $\mathbf{L}' \leftarrow$ UNSPLITTABLE-BLOCKING-FLOW$(\mathbf{L}(D_1), a, \widetilde{w}_i)$;
  if $M^+(\mathbf{L}') = \emptyset$ **then**
    $D_2 \leftarrow \mathsf{ExtractFirst}(\mathcal{L})$;
    $\mathsf{Up}(D_1) \leftarrow \mathsf{Up}(D_2)$;
    $D_1 \leftarrow D_1 \cup D_2$; $\mathsf{Insert}(D_1, \mathcal{L})$;
  else if $M^-(\mathbf{L}') = \emptyset$ **and** $M^+(\mathbf{L}') \neq \emptyset$ **then**
    $\mathsf{Low}(D_1) \leftarrow a$ and output: "links in D_1 are *done in this sweep*";
  else if $M^-(\mathbf{L}') \neq \emptyset$ **and** $M^+(\mathbf{L}') \neq \emptyset$ **then**
    split D_1 (according to Lemma 4.23 (3)) into sets D_1' and $\overline{D_1'}$;
    $\mathsf{Low}(D_1') \leftarrow a$ and output: "links in D_1' are *done in this sweep*";
    $D_2 \leftarrow \mathsf{ExtractFirst}(\mathcal{L})$;
    $\mathsf{Up}(\overline{D_1'}) \leftarrow \mathsf{Up}(D_2)$;
    $D_1 \leftarrow \overline{D_1'} \cup D_2$; $\mathsf{Insert}(D_1, \mathcal{L})$;
  end if
end while
▷ Different handling of last set
$D_1 \leftarrow \mathsf{ExtractFirst}(\mathcal{L})$;
if $a = \mathsf{Up}(D_1) - \widetilde{w}_i$ **then**
  RECURSIVEUBF$(D_1, \mathbf{L}(D_1), [\mathsf{Low}(D_1), a], \widetilde{w}_i)$ and output: "links in D_1 are *done in stage i*";
else
  $\mathbf{L}' \leftarrow$ UNSPLITTABLE-BLOCKING-FLOW$(\mathbf{L}(D_1), a, \widetilde{w}_i)$;
  if $M^-(\mathbf{L}') = \emptyset$ **then**
    $\mathsf{Low}(D_1) \leftarrow a$ and output: "links in D_1 are *done in this sweep*";
  else if $M^-(\mathbf{L}') \neq \emptyset$ **and** $M^+(\mathbf{L}') = \emptyset$ **then**
    $\mathsf{Up}(D_1) \leftarrow a + \widetilde{w}_i$;
    RECURSIVEUBF$(D_1, \mathbf{L}'(D_1), [\mathsf{Low}(D_1), a], \widetilde{w}_i)$ and output: "links in D_1 are *done in stage i*";
  else if $M^-(\mathbf{L}') \neq \emptyset$ **and** $M^+(\mathbf{L}') \neq \emptyset$ **then**
    split D_1 (according to Lemma 4.23 (3)) into sets D_1' and $\overline{D_1'}$;
    $\mathsf{Low}(D_1') \leftarrow a$ and output: "links in D_1' are *done in this sweep*";
    $\mathsf{Up}(\overline{D_1'}) \leftarrow a + \widetilde{w}_i$;
    RECURSIVEUBF$(\overline{D_1'}, \mathbf{L}'(\overline{D_1'}), [\mathsf{Low}(\overline{D_1'}), a], \widetilde{w}_i)$ and output: "links in $\overline{D_1'}$ are *done in stage i*";
  end if
end if

Fig. 4.7. SWEEP over the sets of active links

3. $M^-(\mathbf{L}') \neq \emptyset$ and $M^+(\mathbf{L}') \neq \emptyset$: In this case, we split D_1 according to condition (3) from Lemma 4.23 into sets D_1' and $\overline{D_1'}$. Condition (3c) implies, that no player that is assigned to a link in D_1' can be assigned to a link in $\overline{D_1'}$. Since the load on each link in D_1' is at least a, we can set $\mathsf{Low}(D_1') \leftarrow a$. The load of each link in $\overline{D_1'}$ is at most $a + \widetilde{w}_i$. Thus, since the upper bound of the next element, say D_2, in $\mathcal{L}$ is $\mathsf{Up}(D_2) > a + \widetilde{w}_i$, we again can extract D_2 from $\mathcal{L}$, set $\mathsf{Up}(\overline{D_1'}) \leftarrow \mathsf{Up}(D_2)$, merge $\overline{D_1'}$ and D_2, and insert it in $\mathcal{L}$. We are done with D_1' during this execution of SWEEP.

So, in each case, the number of sets in list $\mathcal{L}$ is decreased by 1. Now, we consider the case that there is only one set, say D_1, in $\mathcal{L}$. This case has to be handled differently.

If $a = \mathsf{Up}(D_1) - \widetilde{w}_i$, then we simply apply RecursiveUBF to the sub-instance defined by D_1 in the interval $[\mathsf{Low}(D_1), a]$ with traffic size $\widetilde{w}_i$. Otherwise, we apply Unsplittable-Blocking-Flow to the sub-instance defined by the set D_1. Unsplittable-Blocking-Flow$(\mathbf{L}(D_1), a, \widetilde{w}_i)$ returns a pure strategy profile $\mathbf{L}'$ where one of the following conditions holds.

1. $M^-(\mathbf{L}') = \emptyset$: Here, we set $\mathsf{Low}(D_1) \leftarrow a$.
2. $M^-(\mathbf{L}') \neq \emptyset$ and $M^+(\mathbf{L}') = \emptyset$: In this case, we set $\mathsf{Up}(D_1) \leftarrow a + \widetilde{w}_i$ and apply RecursiveUBF to the sub-instance defined by D_1 in the interval $[\mathsf{Low}(D_1), a]$ with traffic size $\widetilde{w}_i$.
3. $M^-(\mathbf{L}') \neq \emptyset$ and $M^+(\mathbf{L}') \neq \emptyset$: Here, we split D_1 according to condition (3) from Lemma 4.23 into sets D_1' and $\overline{D_1'}$. For D_1' we set $\mathsf{Low}(D_1') \leftarrow a$ and for $\overline{D_1'}$ we set $\mathsf{Up}(\overline{D_1'}) \leftarrow a + \widetilde{w}_i$ and we apply RecursiveUBF to the sub-instance defined by $\overline{D_1'}$ in the interval $[\mathsf{Low}(\overline{D_1'}), a]$ with traffic size $\widetilde{w}_i$.

After each sweep, by renumbering the partitions, we get a new pure strategy profile that again has the same structure as in Definition 4.26. This completes the description of Sweep. Gairing et al. [41] proved:

Lemma 4.28 (Gairing et al. [41]). *The sweep property holds at the beginning of each execution of* Sweep. *Moreover, in each execution, either a non-empty set of links is added to the set of active links, or some non-empty set of links is done in the current stage.*

Lemma 4.29 (Gairing et al. [41]). *After stage i, every player u with traffic $w_u \geq \widetilde{w}_i$ is satisfied.*

Theorem 4.30 (Gairing et al. [41]). *Consider the model of arbitrary players with restricted strategy sets and identical links. Given an instance $\langle \mathbf{w}, m \rangle$ and an associated pure strategy profile $\mathbf{L}$,* Nashify-Restricted$(\mathbf{L})$ *computes a Nash equilibrium with non-increased makespan social cost in polynomial time.*

Remark 4.31. The algorithm Unsplittable-Blocking-Flow has been proved useful also for the problem of scheduling unrelated parallel machines with the objective to minimize makespan. For this problem, we were able to provide a combinatorial 2-approximation algorithm [51], which is simpler and faster than previously known algorithms. For the approximation algorithm, the procedure Unsplittable-Blocking-Flow is an essential element.

4.4.2 Price of Anarchy

In this section we present a comprehensive collection of bounds on the pure price of anarchy for the model of restricted strategy sets and makespan social cost. Independently of our work, Awerbuch et al. [6] also have studied makespan social cost for the model of restricted strategy sets. They focused on the case of arbitrary players and identical links, for which they proved that the pure price of anarchy is $\Theta(\frac{\log m}{\log\log m})$. Awerbuch et al. [6] also consider mixed Nash equilibria. Here, they showed that the price of anarchy is $\Theta(\frac{\log m}{\log\log\log m})$.

We structure our results on the pure price of anarchy in this model as follows. In Section 4.4.2.1 we prove a lower bound of $\Gamma^{-1}(m)-2$ that holds for the case of identical players and identical links (Theorem 4.4.2.1). This lower bound serves also as a lower bound for the more general cases of arbitrary players or related links (or of both). Section 4.4.2.2 shows an upper bound of $\Gamma^{-1}(n)+1$ for the case of identical players and related links (Theorem 4.33). Section 4.4.2.3 states that for arbitrary players and identical links the pure price of anarchy is upper bounded by $\Theta(\frac{\log m}{\log\log m})$ (Theorem 4.35). We stress that Theorem 4.35 is the only result that can be found in (or even follows from) Awerbuch et al. [6]. Section 4.4.2.4 studies the general model of arbitrary players and related links. Here, we show that the pure price of anarchy lies in between $m-1$ and m (Theorem 4.37 and Theorem 4.38). For our upper bounds in Section 4.4.2.2 and Section 4.4.2.3, we use similar techniques as in [23].

4.4.2.1 Identical Players and Identical Links

We start by proving a lower bound on the pure price of anarchy. This lower bound holds for the model of identical players with restricted strategy sets and identical links.

Theorem 4.32. *Consider the model of identical players with restricted strategy sets and identical links. Then,*

$$\mathsf{pPoA}_{\mathsf{MSP}} > \Gamma^{-1}(m) - 2 = \Omega\left(\frac{\log m}{\log\log m}\right).$$

Proof. Consider an instance $\langle n, m\rangle$ with n identical players with restricted strategy sets and m identical links. We construct the strategy sets of the players as follows. Fix some sufficiently large integer p (to be determined later).

- Partition the set of links into $p+1$ disjoint subsets $\mathcal{M}_0, \mathcal{M}_1, \ldots, \mathcal{M}_p$ with:
  - $|\mathcal{M}_0| = 1$.
  - For each integer l, where $1 \leq l \leq p$, $|\mathcal{M}_l| = (p-1)\cdot\prod_{j\in[l-1]}(p-j)$.

  Note that since $|\mathcal{M}_0| \leq |\mathcal{M}_1| < \ldots < |\mathcal{M}_p|$ the partition implies that $m < (p+1)\cdot|\mathcal{M}_p| = (p+1)(p-1)(p-1)! < (p+1)! = \Gamma(p+2)$. So, $p > \Gamma^{-1}(m)-2$.
- Partition the set of players into p disjoint subsets $\mathcal{U}_0, \mathcal{U}_1, \ldots, \mathcal{U}_{p-1}$ with:
  - For each integer k, where $0 \leq k \leq p-1$, $|\mathcal{U}_k| = (p-k)\cdot|\mathcal{M}_k|$.
  - The strategy set of each player in $\mathcal{U}_k$ is $\mathcal{M}_k \cup \mathcal{M}_{k+1}$.

We now construct a pure Nash equilibrium $\mathbf{L}$ and an optimal strategy profile $\mathbf{Q}$ such that $\mathsf{SC}_{\mathsf{MSP}}(n, m, \mathbf{L}) = p$ and $\mathsf{SC}_{\mathsf{MSP}}(n, m, \mathbf{Q}) = 1$.

- Construct a pure strategy profile $\mathbf{L}$ as follows.

– All p players from the set $\mathcal{U}_0$ are assigned to the single link in $\mathcal{M}_0$.
– For each integer k, where $1 \leq k \leq p-1$, $p-k$ players from $\mathcal{U}_k$ are assigned to each link in $\mathcal{M}_k$. (Note that no player is assigned to any link in $\mathcal{M}_p$.)

By the construction of $\mathbf{L}$, the latency on each link in the set $\mathcal{M}_l$, where $0 \leq l \leq p$, is $p-l$. Thus, for each integer l, where $0 \leq l \leq p-1$, no player assigned to a link in the set $\mathcal{M}_l$ can decrease its private cost by switching either to a different link from the set $\mathcal{M}_l$ or to a link from the set $\mathcal{M}_{l+1}$. So, all players are satisfied in $\mathbf{L}$ and $\mathbf{L}$ is a Nash equilibrium with

$$\begin{aligned}\mathsf{SC}_{\mathsf{MSP}}(n, m, \mathbf{L}) &= \max_{j \in [m]} \Lambda_j(\mathbf{L}) \\ &= \max_{0 \leq l \leq p} (p-l) \\ &= p\,.\end{aligned}$$

- Note that $|\mathcal{M}_0| + |\mathcal{M}_1| = p$ and $|\mathcal{U}_0| = p$. Note also that for each integer k, $1 \leq k \leq p-1$,

$$\begin{aligned}|\mathcal{U}_k| &= (p-k) \cdot |\mathcal{M}_k| \\ &= (p-k)\,(p-1) \cdot \prod_{j \in [k-1]} (p-j) \\ &= (p-1) \cdot \prod_{j \in [k]} (p-j) \\ &= |\mathcal{M}_{k+1}|\,.\end{aligned}$$

So, it is possible to assign each player in $\mathcal{U}_0$ to a distinct link in $\mathcal{M}_0 \cup \mathcal{M}_1$, and to assign each player in $\mathcal{U}_k$, where $1 \leq k \leq p-1$, to a distinct link in $\mathcal{M}_{k+1}$. Call $\mathbf{Q}$ the resulting pure strategy profile. Then, $\mathsf{SC}_{\mathsf{MSP}}(n, m, \mathbf{Q}) = 1$ and $\mathbf{Q}$ is optimal. So, $\mathsf{OPT}_{\mathsf{MSP}}(n, m) = 1$.

It follows that

$$\begin{aligned}\mathsf{pPoA}_{\mathsf{MSP}} &\geq \frac{\mathsf{SC}_{\mathsf{MSP}}(n, m, \mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(n, m)} \\ &= p \\ &> \Gamma^{-1}(m) - 2 \\ &= \Omega\left(\frac{\log m}{\log\log m}\right)\,,\end{aligned}$$

as needed. ∎

Theorem 4.32 implies that $\Omega\left(\frac{\log m}{\log\log m}\right)$ is a lower bound on the pure price of anarchy for the more general cases of arbitrary players or related links (or of both).

4.4.2.2 Identical Players and Related Links

We proceed with an upper bound on the pure price of anarchy for the model of identical players with restricted strategy sets and related links.

Theorem 4.33. *Consider the model of identical players with restricted strategy sets and related links. Then,*

$$\mathsf{pPoA}_{\mathsf{MSP}} \leq \Gamma^{-1}(n) + 1 = O\left(\frac{\log n}{\log \log n}\right).$$

Proof. Consider any arbitrary instance $\langle n, \mathbf{c} \rangle$ with an associated pure Nash equilibrium $\mathbf{L}$ such that

$$k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \leq \mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}, \mathbf{L}) < (k+1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$$

for some integer $k \in \mathbb{N}$, and an optimal strategy profile $\mathbf{Q}$. To prove an upper bound on the price of anarchy, it suffices to prove an upper bound on $k+1$. To do so, we will prove a lower bound (as a function of k) on the number of players that are necessary for such a Nash equilibrium $\mathbf{L}$. We will then use this lower bound to prove an upper bound of $O\left(\frac{\log n}{\log \log n}\right)$ on $k+1$. We continue with the details of the formal proof.

Consider now a link $j \in [m]$ with $c_j < \frac{1}{\mathsf{OPT}_{\mathsf{MSP}}(n,\mathbf{c})}$. Note that in the optimal strategy profile $\mathbf{Q}$, no player is assigned to link j (since otherwise $\frac{1}{c_j} \leq \Lambda_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$, or $c_j \geq \frac{1}{\mathsf{OPT}_{\mathsf{MSP}}(n,\mathbf{c})}$). If, in addition, $\Lambda_j(\mathbf{L}) < \mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}, \mathbf{L})$, then link j can be eliminated (together with all players assigned to it in $\mathbf{L}$) with no change to $\mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}, \mathbf{L})$ and no increase to $\mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$. So, assume, without loss of generality, that for each link $j \in [m]$, either $c_j \geq \frac{1}{\mathsf{OPT}_{\mathsf{MSP}}(n,\mathbf{c})}$ or $\Lambda_j(\mathbf{L}) = \mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}, \mathbf{L})$.

Define $\mathcal{M}_0$ as the set of links $j \in [m]$ with latency

$$\Lambda_j(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}).$$

Clearly, $\mathcal{M}_0 \neq \emptyset$. By definition of latency, this implies that

$$\sum_{j \in \mathcal{M}_0} \delta_j(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j .$$

We prove an inductive claim:

Lemma 4.34. *For each $l \in [k-1]$, there is a set of links $\mathcal{M}_l$ with $\mathcal{M}_l \cap (\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_{l-1}) = \emptyset$ such that:*

(1) $\sum_{j \in \mathcal{M}_l} c_j \geq (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \sum_{j \in \mathcal{M}_0} c_j$.
(2) *For each link $j \in \mathcal{M}_l$, $\Lambda_j(\mathbf{L}) \geq (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$.*
(3) $\sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{L}) \geq (k-1) \cdot \prod_{j \in [l]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j$.
(4) *There are at least $(k-1) \cdot \prod_{j \in [l]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j$ players assigned by $\mathbf{L}$ to links in $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_l$ whose strategy sets include links outside $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_l$.*

Proof. By (strong) induction on l. For the sake of shortening the proof, we merge the proof for the basis case (where $l = 1$) into the proof for the induction step;

thus, the case $l = 1$ will be treated separately (where needed) along the proof of the induction step.

Assume inductively that for some integer $l \geq 1$, the claim holds for all integers not exceeding $(l-1)$. Notice that if $l = 1$, the induction hypothesis is empty. We will prove the claim for l.

Assume first that $l = 1$. Recall that

$$\sum_{j \in \mathcal{M}_0} \delta_j(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j .$$

In the optimal strategy profile $\mathbf{Q}$, $\Lambda_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$ for each link $j \in [m]$. By definition of latency, this implies that $\sum_{j \in \mathcal{M}_0} \delta_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j$. It follows that there are at least

$$\begin{aligned} & k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j - \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\ & = (k-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \end{aligned}$$

excess players assigned by $\mathbf{L}$ to links in $\mathcal{M}_0$ whose strategy sets include links outside $\mathcal{M}_0$.

Assume now that $l > 1$. By induction hypothesis (condition (4)), there are at least $(k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j$ excess players assigned by $\mathbf{L}$ to links in $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_{l-1}$ whose strategy sets include links outside $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_{l-1}$.

Define $\mathcal{M}_l$ as the set of all links outside $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_{l-1}$ that are included in the strategy sets of such excess players; so, $\mathcal{M}_l \cap (\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_{l-1}) = \emptyset$.

Clearly, in $\mathbf{Q}$, all these excess players are assigned to links in $\mathcal{M}_l$, so that

$$\sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{Q}) \geq (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j .$$

We now prove the four claimed properties for the set $\mathcal{M}_l$.

- Clearly,

$$\begin{aligned} \sum_{j \in \mathcal{M}_l} c_j &= \sum_{j \in \mathcal{M}_l} \frac{\delta_j(\mathbf{Q})}{\Lambda_j(\mathbf{Q})} \\ &\geq \frac{\sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{Q})}{\mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})} \\ &\geq \frac{(k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j}{\mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})} \\ &= (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \sum_{j \in \mathcal{M}_0} c_j , \end{aligned}$$

  which proves (1).

- To prove (2), consider any link $j \in \mathcal{M}_l$. Since $j \notin \mathcal{M}_0$, it follows that $\Lambda_j(\mathbf{L}) < k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$. Since $\mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$, this implies that $\Lambda_j(\mathbf{L}) < \mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c})$. Therefore, $c_j \geq \frac{1}{\mathsf{OPT}_{\mathsf{MSP}}(n,\mathbf{c})}$.
  If $l = 1$, there is some link $j' \in \mathcal{M}_0$ to which $\mathbf{L}$ assigns some excess player. So, assume $l > 1$. Recall that in the optimal strategy profile $\mathbf{Q}$, $\Lambda_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$ for each link $j \in \mathcal{M}_l$. By definition of latency, this implies that $\sum_{j \in \mathcal{M}_{l-1}} \delta_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_{l-1}} c_j$. By induction hypothesis (condition (3)),
  $$\sum_{j \in \mathcal{M}_{l-1}} \delta_j(\mathbf{L}) \geq (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j .$$
  It follows that there is some excess player assigned to some link $j' \in \mathcal{M}_{l-1}$. Since $\mathbf{L}$ is a Nash equilibrium, for each link $j \in \mathcal{M}_l$,
  $$\begin{aligned} \Lambda_{j'}(\mathbf{L}) &\leq \Lambda_j(\mathbf{L}) + \frac{1}{c_j} \\ &\leq \Lambda_j(\mathbf{L}) + \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) . \end{aligned}$$
  Assume first that $l = 1$. By definition of the set $\mathcal{M}_0$,
  $$\Lambda_{j'}(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) .$$
  It follows that
  $$\Lambda_j(\mathbf{L}) \geq (k-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}),$$
  and the proof of (2) for the basis case is now complete.
  So, assume $l > 1$. By induction hypothesis (condition (2)),
  $$\begin{aligned} \Lambda_{j'}(\mathbf{L}) &\geq (k-(l-1)) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \\ &= (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) + \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) . \end{aligned}$$
  It follows that
  $$\Lambda_j(\mathbf{L}) \geq (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}),$$
  and the proof of (2) is now complete.
- To prove (3), we use (2) and (1) to derive that
  $$\begin{aligned} \sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{L}) &= \sum_{j \in \mathcal{M}_l} \Lambda_j(\mathbf{L}) \cdot c_j \\ &\geq (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j \\ &\geq (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \sum_{j \in \mathcal{M}_0} c_j \\ &= (k-1) \cdot \prod_{j \in [l]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j , \end{aligned}$$
  as needed for proving (3).

- Recall first that in the optimal strategy profile $\mathbf{Q}$, $\Lambda_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$ for each link $j \in [m]$. By definition of latency, this implies that $\sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j$.
  Clearly, the number of players assigned by $\mathbf{L}$ to links in $\mathcal{M}_0 \cup \ldots \cup \mathcal{M}_l$ whose strategy sets include links outside $\mathcal{M}_0 \cup \ldots \mathcal{M}_l$ is at least

$$\begin{aligned}
&\sum_{r \in [0,l]} \sum_{j \in \mathcal{M}_r} (\delta_j(\mathbf{L}) - \delta_j(\mathbf{Q})) \\
&\quad = \sum_{r \in [0,l-1]} \sum_{j \in \mathcal{M}_r} (\delta_j(\mathbf{L}) - \delta_j(\mathbf{Q})) + \sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{L}) - \sum_{j \in \mathcal{M}_l} \delta_j(\mathbf{Q}) \\
&\quad \geq (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\qquad + (k-l) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j - \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j \\
&\quad = (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\qquad + (k-l-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_l} c_j \\
&\quad \geq (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\qquad + (k-l-1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\quad = (1 + (k-l-1)) \cdot (k-1) \cdot \prod_{j \in [l-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\quad = (k-1) \cdot \prod_{j \in [l]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \, ,
\end{aligned}$$

  as needed for proving (4).

The proof of the inductive claim is now complete. ■

We now prove an upper bound on $k+1$. Fix any link $j \in \mathcal{M}_0$. Clearly, $\Lambda_j(\mathbf{L}) \leq \mathsf{SC}_{\mathsf{MSP}}(n, \mathbf{c}, \mathbf{L}) < (k+1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$. Recall that by definition of $\mathcal{M}_0$, $\Lambda_j(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) > 0$. This implies that $\Lambda_j(\mathbf{L}) \geq \frac{1}{c_j}$. It follows that $\frac{1}{c_j} < (k+1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c})$. This implies that

$$\mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j > \frac{1}{k+1} \, .$$

Assume, without loss of generality, that $k \geq 3$ (otherwise $k+1 \in O(1)$). Then, by Lemma 4.34 (condition (3)),

$$\begin{aligned}
n &\geq \sum_{j \in \mathcal{M}_{k-1}} \delta_j(\mathbf{L}) + \sum_{j \in \mathcal{M}_{k-2}} \delta_j(\mathbf{L}) \\
&\geq (k-1) \cdot \prod_{j \in [k-1]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&\quad + (k-1) \cdot \prod_{j \in [k-2]} (k-j) \cdot \mathsf{OPT}_{\mathsf{MSP}}(n, \mathbf{c}) \cdot \sum_{j \in \mathcal{M}_0} c_j \\
&> 2 \cdot (k-1) \cdot (k-1)! \cdot \frac{1}{k+1} \\
&\geq (k-1)! \\
&= \Gamma(k)\,.
\end{aligned}$$

Hence

$$\begin{aligned}
k+1 &< \Gamma^{-1}(n) + 1 \\
&= O\left(\frac{\log n}{\log\log n}\right) ,
\end{aligned}$$

as needed. ■

We remark that Theorems 4.32 and 4.33 leave a gap between our bounds on the pure price of anarchy for the case of identical players. Closing this gap remains an interesting open problem.

4.4.2.3 Arbitrary Players and Identical Links

With a similar proof as in Theorem 4.33, we can prove an upper bound on the pure price of anarchy for the model of arbitrary players with restricted strategy sets and identical links. This upper bound matches asymptotically the lower bound shown in Theorem 4.32.

Theorem 4.35. *Consider the model of arbitrary players with restricted strategy sets and identical links. Then,*

$$\mathsf{pPoA}_{\mathsf{MSP}} = \Gamma^{-1}(m) = O\left(\frac{\log m}{\log\log m}\right) .$$

Theorems 4.33 and 4.35 together imply:

Theorem 4.36. *Consider the model of identical players with restricted strategy sets and identical links. Then,*

$$\mathsf{pPoA}_{\mathsf{MSP}} = O\left(\frac{\log \min\{m,n\}}{\log\log \min\{m,n\}}\right) .$$

We remark that in the interesting cases where $n \geq m$, Theorems 4.32 and 4.36 provide asymptotically tight bounds on the pure price of anarchy for the case of identical players and identical links.

4.4.2.4 Arbitrary Players and Related Links

We now turn to the model of arbitrary players with restricted strategy sets and related links. For this model, we provide almost matching upper and lower bounds on the pure price of anarchy. We first prove the lower bound:

Theorem 4.37. *Consider the model of arbitrary players with restricted strategy sets and related links. Then,* $\mathsf{pPoA} \geq m-1$.

Proof. Consider an instance $\langle \mathbf{w}, \mathbf{c} \rangle$ as follows:

- For each link $j \in [m]$, the capacity c_j is

$$c_j = \frac{(m-1)!}{(j-1)!}\,.$$

- There are $n = m-1$ players; the weight of player $i \in [m-1]$ is $w_i = c_i$.

Moreover, assume that for each player $i \in [m-1]$, the strategy set S_i is $S_i = \{i, i+1\}$.

- Construct a pure strategy profile $\mathbf{L}$ as follows:
  Each player $i \in [m-1]$ is assigned to link $i+1$.
  We will argue that all players are satisfied in $\mathbf{L}$.
  - On the one hand, the private cost of each player $i \in [m-1] \setminus \{1\}$ is

$$\begin{aligned} \mathsf{PC}_i(\mathbf{L}) &= \Lambda_{i+1}(\mathbf{L}) \\ &= \frac{w_i}{c_{i+1}} \\ &= i\,. \end{aligned}$$

    On the other hand, moving to the other link i in its strategy set would lead to latency

$$\begin{aligned} \frac{\delta_i + w_i}{c_i} &= \frac{w_{i-1} + w_i}{c_i} \\ &= \frac{c_{i-1} + c_i}{c_i} \\ &= (i-1) + 1 \\ &= i\,. \end{aligned}$$

    It follows that player $i \in [m-1] \setminus \{1\}$ is satisfied in $\mathbf{L}$.
  - Consider now player 1. Since $c_1 = c_2$ and there are no players assigned to link 1, player 1 cannot decrease its private cost by switching from link 2 to link 1. So, player 1 is also satisfied in $\mathbf{L}$.

  It follows that $\mathbf{L}$ is a Nash equilibrium. Clearly,

$$\begin{aligned} \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) &= \max_{j \in [m]} \Lambda_j(\mathbf{L}) \\ &= m-1\,. \end{aligned}$$

- Construct now a pure strategy profile $\mathbf{Q}$ as follows:
  Each player $i \in [m-1]$ is assigned to link i.
  Clearly, for each link $j \in [m-1]$, $\Lambda_j(\mathbf{L}) = \frac{w_j}{c_j} = 1$ and $\Lambda_m(\mathbf{L}) = 0$. So, $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{Q}) = 1$. Thus, $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \leq 1$.

It follows that

$$\begin{aligned}\mathsf{pPoA}_{\mathsf{MSP}} &\geq \frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L})}{\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})} \\ &= m - 1\,,\end{aligned}$$

as needed. ■

We now prove the upper bound:

Theorem 4.38. *Consider the model of arbitrary players with restricted strategy sets and related links. Then,* $\mathsf{pPoA} < m$.

Proof. Consider any arbitrary instance $\langle \mathbf{w}, \mathbf{c} \rangle$ with an associated Nash equilibrium $\mathbf{L}$ such that

$$k \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) \leq \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) < (k+1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$$

for some integer $k \in \mathbb{N}$, and an optimal strategy profile $\mathbf{Q}$. To prove an upper bound on the price of anarchy, it suffices to prove an upper bound on $k+1$. To do so, we will prove a lower bound (as a function of k) on the number of links that are necessary for such a Nash equilibrium $\mathbf{L}$. We will then use this lower bound to prove an upper bound of m on $k+1$. We continue with the details of the formal proof.

We prove an inductive claim:

Lemma 4.39. *For each integer $i \in [k]$, there exists a distinct link $l_i \in [m]$ with latency $\Lambda_{l_i}(\mathbf{L}) \geq (k - i + 1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$.*

Proof. By (strong) induction on i. For the basis case, let $i = 1$. Since $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$, there is a link $l_1 \in [m]$ with latency $\Lambda_{l_1}(\mathbf{L}) \geq k \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})$, as needed.

Assume inductively that the claim holds for all integers not exceeding $(i-1)$ where $i \geq 2$. We will prove the claim for i. By induction hypothesis, there exist $i-1$ distinct links $l_1, \ldots, l_{i-1}$ with

$$\Lambda_{l_j}(\mathbf{L}) \geq (k - j + 1) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})\,,$$

for each integer $j \in [i-1]$. Since $j \leq i-1$ and $i \leq k$, it follows that $j \leq k-1$. So, $k - j + 1 \geq 2$. It follows that for each integer $j \in [i-1]$,

$$\Lambda_{l_j}(\mathbf{L}) > \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c})\,.$$

Since $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{w}, \mathbf{c}, \mathbf{Q}) \geq \Lambda_{l_j}(\mathbf{Q})$ for each integer $j \in [i-1]$, it follows that for each integer $j \in [i-1]$, $\Lambda_{l_j}(\mathbf{L}) > \Lambda_{l_j}(\mathbf{Q})$. So, $\sum_{j \in [i-1]} \Lambda_{l_j}(\mathbf{L}) >$

$\sum_{j\in[i-1]} \Lambda_{l_j}(\mathbf{Q})$. It follows that there is some player i_0 assigned by $\mathbf{L}$ to some link in the set $\{l_1,\ldots,l_{i-1}\}$ that is assigned by $\mathbf{Q}$ to some link $l_i \notin \{l_1,\ldots,l_{i-1}\}$ (otherwise, $\sum_{j\in[i-1]} \Lambda_{l_j}(\mathbf{Q}) \geq \sum_{j\in[i-1]} \Lambda_{l_j}(\mathbf{L})$). Thus, l_i is an allowed link for player i_0.

Since $\mathbf{L}$ is a Nash equilibrium, player i_0 has no incentive to switch from its link l_j, where $j \in [i-1]$, to link l_i. Since player i_0 is assigned to link l_i in $\mathbf{Q}$, the additional latency on link l_i in $\mathbf{L}$ due to player i_0 switching to link l_i is at most the latency on link l_i in $\mathbf{Q}$; since $\mathbf{Q}$ is optimal, this additional latency is at most $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$. It follows that

$$\Lambda_{l_j}(\mathbf{L}) \leq \Lambda_{l_i}(\mathbf{L}) + \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\,.$$

By induction hypothesis,

$$\begin{aligned}\Lambda_{l_j}(\mathbf{L}) &\geq (k-j+1)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\\ &\geq (k-(i-1)+1)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\\ &= (k-i+1)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c}) + \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\,.\end{aligned}$$

It follows that

$$\Lambda_{l_i}(\mathbf{L}) \geq (k-i+1)\cdot \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})\,.$$

The proof of the inductive claim is now complete. ■

Lemma 4.39 implies that for $\mathbf{L}$, there are k distinct links with latency larger than $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$. Since $\sum_{j\in[m]} \Lambda_j(\mathbf{L}) = \sum_{j\in[m]} \Lambda_j(\mathbf{Q})$ and $\Lambda_j(\mathbf{Q}) \leq \mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$ for each link $j \in [m]$, it follows that there is some other link with latency smaller than $\mathsf{OPT}_{\mathsf{MSP}}(\mathbf{w},\mathbf{c})$. So, $k \leq m-1$ or $k+1 \leq m$, as needed. ■

4.5 Polynomial Social Cost

We now come back to the case of unrestricted strategy sets and study *polynomial social cost* for routing games on parallel links. Throughout this section, we restrict to the case of *identical links*. Recall, that polynomial social cost is defined with the help of a certain polynomial cost function $\pi_d(\lambda)$ of degree d. Using a different definition of social cost does not alter the definition of private cost nor the set of (pure) Nash equilibria. Thus, the results on sequences of selfish steps from Section 4.2.1 also apply to this model. Furthermore, Lücking [70] showed that selfish steps do not increase polynomial social cost. This implies that NASHIFY-IDENTICAL from Section 4.2.2 can be used to convert a given strategy profile into a pure Nash equilibrium with non-increased polynomial social cost.

In this section, we are interested in the price of anarchy for polynomial social cost. We start by proving a simple fact (Lemma 4.40), that will be instrumental for reducing the polynomial price of anarchy for arbitrary polynomials (with non-negative coefficients) to the monomial price of anarchy. This result holds for arbitrary players. We then focus on the case of identical players. Here, we show that the monomial social cost of the fully mixed Nash equilibrium can be expressed as a combinatorial sum of Stirling numbers of the second kind (Corollary 4.41). With the help of Corollary 4.41, we show that fully mixed Nash equilibria maximize polynomial social cost, for the case of identical players and two identical links (Theorem 4.42). Afterwards, we consider the case of identical players and arbitrary many identical links. Here, we show that fully mixed Nash equilibria maximize polynomial social cost up to a factor $\left(1+\frac{1}{n-1}\right)^d$ (Theorem 4.44). Recently, it was shown that this factor is not necessary (Theorem 4.47).

Equipped with these results, we show that for the model of identical players and identical links, the price of anarchy is upper bounded by B_d. Recall that B_d is the *Bell number* of order d. Our analysis first shows that B_d is an upper bound on the price of anarchy, if the polynomial cost function is the d'th power (Theorem 4.48). As a corollary we get that the same upper bound also holds for general polynomial cost functions (Corollary 4.49). We show in Theorem 4.50 and Corollary 4.51 that, for the special case of 2 links, both upper bounds reduce from B_d to $2^{d-2}\left(1+\left(\frac{1}{n}\right)^{d-1}\right)$.

We start by proving a simple fact that holds for arbitrary players.

Lemma 4.40 (From Polynomials to Monomials). *Consider the model of arbitrary players and identical links. Fix an instance $\langle \mathbf{w}, m\rangle$ with an associated Nash equilibrium $\mathbf{P}$. Then,*

$$\frac{\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{P})}{\mathsf{OPT}_{\pi_d(\lambda)}(\mathbf{w},m)} \le \max_{t\in[d]}\left\{\frac{\mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{P})}{\mathsf{OPT}_{\lambda^t}(\mathbf{w},m)}\right\}.$$

Proof. Our proof will use the expression of polynomial social cost as a linear combination of monomial social costs (see Section 3.1.5.2). The proof will manipulate sums of fractions while relying on the non-negativeness of the coefficients in the latency cost function. We continue with the details of the formal proof.

Let $\mathbf{Q}$ be an optimal strategy profile for the instance $\langle \mathbf{w}, m\rangle$. Then,

$$\begin{aligned}\frac{\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{P})}{\mathsf{OPT}_{\pi_d(\lambda)}(\mathbf{w},m)} &= \frac{\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{P})}{\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{Q})}\\ &= \frac{a_0+\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{P})}{a_0+\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{Q})}.\end{aligned}$$

Since $\mathbf{Q}$ is an optimal strategy profile, $\mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{P}) \ge \mathsf{SC}_{\pi_d(\lambda)}(\mathbf{w},m,\mathbf{Q})$. Since the coefficients a_t are non-negative for all $t\in[d]$, it follows that $\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{P}) \ge \sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{Q})$. Since $a_0\ge 0$, this implies that

$$\frac{a_0+\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{P})}{a_0+\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{Q})} \le \frac{\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{P})}{\sum_{t\in[d]} a_t\cdot \mathsf{SC}_{\lambda^t}(\mathbf{w},m,\mathbf{Q})}.$$

Now define, for each $t \in [d]$,

$$K(t, \mathbf{P}, \mathbf{Q}) = \frac{\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})}.$$

Observe, that for each $t \in [d]$, $\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q}) > 0$; so $K(t, \mathbf{P}, \mathbf{Q})$ is well defined. It follows that

$$\begin{aligned}
&\frac{\sum_{t\in[d]} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\sum_{t\in[d]} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\
&\quad = \frac{\sum_{t\in[d]} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q}) \cdot K(t, \mathbf{P}, \mathbf{Q})}{\sum_{t\in[d]} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\
&\quad \leq \frac{\max_{t\in[d]} \{K(t, \mathbf{P}, \mathbf{Q})\} \cdot \sum_{t\in[d]} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})}{\sum_{t\in[d] | a_t > 0} a_t \cdot \mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\
&\quad = \max_{t\in[d]} \{K(t, \mathbf{P}, \mathbf{Q})\} \\
&\quad = \max_{t\in[d]} \left\{ \frac{\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \right\} .
\end{aligned}$$

Since $\mathsf{SC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q}) \geq \mathsf{OPT}_{\lambda^t}(\mathbf{w}, m)$ the claim follows. ■

4.5.1 Identical Players

In the following we restrict to the case of identical players. In this case, by Definition 2.1, polynomial social cost reduces to

$$\begin{aligned}
\mathsf{SC}_{\pi_d(\lambda)}(n, m, \mathbf{P}) &= \sum_{j\in[m]} \sum_{A\subseteq[n]} \left(\prod_{i\in A} p_{ij} \right) \left(\prod_{i\notin A} (1 - p_{ij}) \right) \pi_d(|A|) \\
&= \sum_{j\in[m]} \mathsf{BF}(\langle p_{1j}, \ldots, p_{nj} \rangle, \pi^d(\lambda)) .
\end{aligned}$$

So, polynomial social cost is a sum of binomial functions, one for each link.

4.5.1.1 Fully Mixed Nash Equilibria

We proceed by providing a simple expression for the monomial social cost of the fully mixed Nash equilibrium. Recall that in the case of identical links, all probabilities are identical (and equal to $\frac{1}{m}$) for the fully mixed Nash equilibrium $\mathbf{F}$. Hence, Proposition 2.2 implies now that the monomial social cost of the fully mixed Nash equilibrium $\mathbf{F}$ is a combinatorial sum of Stirling numbers of the second kind.

Corollary 4.41. *Consider the model of identical players and identical links. Fix an instance $\langle n, m \rangle$. Then,*

$$\mathsf{SC}_{\lambda^d}(n, m, \mathbf{F}) = m \sum_{t\in[d]} \left(\frac{1}{m} \right)^t \cdot S(d, t) \cdot n^{\underline{t}} .$$

A lower bound on the monomial optimum for the case of identical players is $\mathsf{OPT}_{\lambda^d}(n,m) \geq m\left(\frac{n}{m}\right)^d$ if $n \geq m$, while $\mathsf{OPT}_{\lambda^d}(n,m) = n$ if $n < m$.

We are now ready to prove that fully mixed Nash equilibria maximize polynomial social cost for the case of identical players and two identical links.

Theorem 4.42. *Consider the model of identical players and two identical links. Fix an instance $\langle n, 2\rangle$ with associated Nash equilibrium $\mathbf{P}$ and fully mixed Nash equilibrium $\mathbf{F}$. Then,*

$$\mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{F}) .$$

Proof. Since polynomial social cost is a linear combination (with non-negative coefficients) of monomial social costs, it suffices to prove the claim for a monomial latency cost function $\pi_d(\lambda) = \lambda^d$.

We partition the set of players $[n]$ into three sets:

$$\mathcal{U}_1 = \{i \in [n] \mid \mathsf{support}_i(\mathbf{P}) = \{1\}\},$$
$$\mathcal{U}_2 = \{i \in [n] \mid \mathsf{support}_i(\mathbf{P}) = \{2\}\},$$
$$\mathcal{U}_{12} = \{i \in [n] \mid \mathsf{support}_i(\mathbf{P}) = \{1, 2\}\}.$$

Without loss of generality, assume that $|\mathcal{U}_1| \leq |\mathcal{U}_2|$. Denote $u = |\mathcal{U}_1|$, $v = |\mathcal{U}_2| - u$ and $r = |\mathcal{U}_{12}|$.

We will treat separately pure and non-pure Nash equilibria. In the second case, we will distinguish between the subcase where there are no pure players assigned to one of the two links (that is, $u = 0$), and the subcase where there are pure players assigned to each of the two links (that is, $u > 0$). The second subcase will be reduced to the first. We now continue with the details of the formal proof. We proceed by case analysis.

1. Assume first that $\mathbf{P}$ is pure. Since $\mathbf{P}$ is a Nash equilibrium, $\Lambda_1(\mathbf{P}) \leq \Lambda_2(\mathbf{P}) + 1$ and $\Lambda_2(\mathbf{P}) \leq \Lambda_1(\mathbf{P}) + 1$. So, $|\Lambda_1(\mathbf{P}) - \Lambda_2(\mathbf{P})| \leq 1$. Note that $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) = (\Lambda_1(\mathbf{P}))^d + (\Lambda_2(\mathbf{P}))^d$ and $\Lambda_1(\mathbf{P}) + \Lambda_2(\mathbf{P}) = n$. Hence, $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P})$ is minimum when $|\Lambda_1(\mathbf{P}) - \Lambda_2(\mathbf{P})| \leq 1$. It follows that $\mathbf{P}$ is an optimal strategy profile, so that $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{F})$, as needed.
2. Assume now that $\mathbf{P}$ is not pure, so that $r > 0$. There are two separate cases.
   a) Assume first that $u = 0$. So, no pure player is assigned to link 1. We first prove that, in this case, $r > 1$. Assume, by way of contradiction, that $r = 1$, and consider the single mixed player i_0. Then, $\lambda_{i_0 1}(\mathbf{P}) = 1$, while $\lambda_{i_0 2}(\mathbf{P}) = |\mathcal{U}_2| + 1$. Since $\mathbf{P}$ is a Nash equilibrium, $\lambda_{i_0 1}(\mathbf{P}) = \lambda_{i_0 2}(\mathbf{P})$. It follows that $|\mathcal{U}_2| = 0$. This implies that $n = r = 1$, a contradiction. It follows that $r > 1$. So, $r \in [n-1] \setminus \{1\}$ mixed players are assigned to both links and $n - r$ pure players are assigned to link 2.
   Consider any arbitrary player $i \in \mathcal{U}_{12}$. Clearly,
   $$\lambda_{i1}(\mathbf{P}) = \Lambda_1(\mathbf{P}) - p_{i1} + 1$$

and

$$\begin{aligned}\lambda_{i2}(\mathbf{P}) &= \Lambda_2(\mathbf{P}) - p_{i2} + 1 \\ &= \Lambda_2(\mathbf{P}) - (1 - p_{i1}) + 1.\end{aligned}$$

Since $\mathbf{P}$ is a Nash equilibrium, $\lambda_{i1}(\mathbf{P}) = \lambda_{i2}(\mathbf{P})$. This implies that

$$p_{i1} = \frac{\Lambda_1(\mathbf{P}) - \Lambda_2(\mathbf{P}) + 1}{2}.$$

So, p_{i1} is independent of i. It follows that each mixed player chooses link 1 with probability $p_1 = p_{i1}$ and link 2 with probability $p_2 = 1 - p_1$. Hence, $\Lambda_1(\mathbf{P}) = rp_1$ and $\Lambda_2(\mathbf{P}) = (n-r) + r(1-p_1)$. We obtain that $p_1 = \frac{rp_1 - ((n-r)+r(1-p_1))+1}{2}$, from which we derive that

$$\begin{aligned}p_1 &= \frac{1}{2} + \frac{n-r}{2(r-1)}\,, \\ p_2 &= \frac{1}{2} - \frac{n-r}{2(r-1)}\,.\end{aligned}$$

Clearly, $\Lambda_1(\mathbf{P}) = rp_1$ and $\Lambda_2(\mathbf{P}) = n - \Lambda_1(\mathbf{P})$. Denote $\alpha = \frac{n}{2}$ and $\beta = \frac{n-r}{2(r-1)}$ to derive that $\Lambda_1(\mathbf{P}) = \alpha + \beta$ and $\Lambda_2(\mathbf{P}) = \alpha - \beta$. Then, the average probabilities on links 1 and 2 are

$$\begin{aligned}\widetilde{p}_1 &= \frac{\Lambda_1(\mathbf{P})}{r} = \frac{\alpha+\beta}{r}, \text{ and} \\ \widetilde{p}_2 &= \frac{\Lambda_2(\mathbf{P})}{n} = \frac{\alpha-\beta}{n},\end{aligned}$$

respectively.

On one hand, Lemma 2.3 and Proposition 2.2 imply that

$$\begin{aligned}\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) &= \mathsf{BF}\left(\langle \underbrace{p_1, \ldots, p_1}_{r \text{ entries}} \rangle, \lambda^d\right) \\ &\quad + \mathsf{BF}\left(\langle \underbrace{p_2, \ldots, p_2}_{r \text{ entries}}, \underbrace{1, \ldots, 1}_{n-r \text{ entries}} \rangle, \lambda^d\right) \\ &\leq \mathsf{BF}\left(\widetilde{p}_1, r, \lambda^d\right) + \mathsf{BF}\left(\widetilde{p}_2, n, \lambda^d\right) \\ &= \sum_{t \in [d]} \left((\widetilde{p}_1)^t \cdot S(d,t) \cdot r^{\underline{t}}\right) + \sum_{t \in [d]} \left((\widetilde{p}_2)^t \cdot S(d,t) \cdot n^{\underline{t}}\right) \\ &= \sum_{t \in [d]} S(d,t) \left((\alpha+\beta)^t \cdot \frac{r^{\underline{t}}}{r^t} + (\alpha-\beta)^t \cdot \frac{n^{\underline{t}}}{n^t}\right)\,.\end{aligned}$$

On the other hand, Lemma 4.41 and Proposition 2.2 imply that

$$\begin{aligned}
\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{F}) &= \mathsf{BF}\left(\langle f_{11}, \ldots, f_{n1}\rangle, \lambda^d\right) + \mathsf{BF}\left(\langle f_{12}, \ldots, f_{n2}\rangle, \lambda^d\right) \\
&= 2\,\mathsf{BF}\left(\frac{1}{2}, n, \lambda^d\right) \\
&= 2\,\mathsf{BF}\left(\frac{\alpha}{n}, n, \lambda^d\right) \\
&= 2 \sum_{t \in [d]} \left(\frac{\alpha}{n}\right)^t \cdot S(d,t) \cdot n^{\underline{t}} \\
&= \sum_{t \in [d]} S(d,t) \cdot \left(2\,\alpha^t \cdot \frac{n^{\underline{t}}}{n^t}\right) .
\end{aligned}$$

So, clearly,

$$\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{F}) - \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) = \sum_{t \in [d]} S(d,t) \cdot \Delta(t)\,,$$

where for each integer $t \in [d]$,

$$\Delta(t) = 2\,\alpha^t \cdot \frac{n^{\underline{t}}}{n^t} - \left((\alpha+\beta)^t \cdot \frac{r^{\underline{t}}}{r^t} + (\alpha-\beta)^t \cdot \frac{n^{\underline{t}}}{n^t}\right) .$$

We prove:

Lemma 4.43. *For each integer $t \geq 1$, $\Delta(t) \geq 0$.*

Proof. By induction on t. For the basis case where $t = 1$, the claim holds since $2\alpha - ((\alpha+\beta) + (\alpha-\beta)) = 0$. Assume inductively that the claim holds for $(t-1)$, where $t \geq 2$. For the induction step, we will prove the claim for t.
Note first that by the definition of α and β and since $r \leq n$,

$$\begin{aligned}
(\alpha+\beta) \cdot \frac{r-(t-1)}{r} &= \frac{(n-1)r}{2(r-1)} \cdot \frac{r-(t-1)}{r} \\
&= \frac{n-1}{2} \cdot \frac{r-(t-1)}{r-1} \\
&\leq \frac{n-1}{2} \cdot \frac{n-(t-1)}{n-1} \\
&= \frac{n}{2} \cdot \frac{n-(t-1)}{n} \\
&= \alpha \cdot \frac{n-(t-1)}{n} .
\end{aligned}$$

We now use this fact to derive that

$$
\begin{aligned}
&(\alpha+\beta)^t \cdot \frac{r^{\underline{t}}}{r^t} + (\alpha-\beta)^t \cdot \frac{n^{\underline{t}}}{n^t} \\
&= (\alpha+\beta) \cdot \frac{r-(t-1)}{r} \cdot (\alpha+\beta)^{t-1} \cdot \frac{r^{\underline{t-1}}}{r^{t-1}} \\
&\quad +(\alpha-\beta) \cdot \frac{n-(t-1)}{n} \cdot (\alpha-\beta)^{t-1} \cdot \frac{n^{\underline{(t-1)}}}{n^{t-1}} \\
&\leq \alpha \cdot \frac{n-(t-1)}{n} \cdot (\alpha+\beta)^{t-1} \cdot \frac{r^{\underline{t-1}}}{r^{t-1}} \\
&\quad +(\alpha-\beta) \cdot \frac{n-(t-1)}{n} \cdot (\alpha-\beta)^{t-1} \cdot \frac{n^{\underline{(t-1)}}}{n^{t-1}} \\
&\leq \alpha \cdot \frac{n-(t-1)}{n} \cdot (\alpha+\beta)^{t-1} \cdot \frac{r^{\underline{t-1}}}{r^{t-1}} \\
&\quad +\alpha \cdot \frac{n-(t-1)}{n} \cdot (\alpha-\beta)^{t-1} \cdot \frac{n^{\underline{(t-1)}}}{n^{t-1}} \\
&\leq \alpha \cdot \frac{n-(t-1)}{n} \cdot \left((\alpha+\beta)^{t-1} \cdot \frac{r^{\underline{t-1}}}{r^{t-1}} + (\alpha-\beta)^{t-1} \cdot \frac{n^{\underline{(t-1)}}}{n^{t-1}} \right) \\
&\leq \alpha \cdot \frac{n-(t-1)}{n} \cdot 2\alpha^{t-1} \cdot \frac{n^{\underline{t-1}}}{n^{t-1}} \\
&= 2\alpha^t \cdot \frac{n^{\underline{t}}}{n^t} \,,
\end{aligned}
$$

where we used the induction hypothesis in the last inequality. This completes the proof of the lemma. ■

Lemma 4.43 implies that $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{F})$. The proof for the case $u = 0$ that $\mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{F})$ is now complete.

b) Assume now that $u > 0$.
Consider the mixed strategy profile $\mathbf{Q}$ for the instance $\langle n, 2\rangle$, which assigns u pure players to each link (with probability 1) and $\widetilde{n} = n - 2u$ mixed players to each link with probability $\frac{1}{2}$. Clearly, $\mathbf{Q}$ is a Nash equilibrium. Note that the average probability for each link is $\frac{u \cdot 1 + (n-2u)\cdot \frac{1}{2}}{n} = \frac{1}{2}$, which is precisely the probability with which each player is assigned to a link in the fully mixed Nash equilibrium $\mathbf{F}$. Since polynomial social cost is a sum of binomial functions and the function $\pi_d(\lambda) = \lambda^d$ is convex, Lemma 2.3 implies that

$$\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{Q}) \leq \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{F}).$$

In the rest, we will prove that $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{Q})$, and this will complete the proof.
Denote $\widetilde{\mathbf{F}}$ the (unique) fully mixed Nash equilibrium associated with the instance $\langle \widetilde{n}, 2\rangle$. On one hand,

$$\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{Q}) = \mathsf{SC}_{(\lambda+u)^d}(\widetilde{n}, 2, \widetilde{\mathbf{F}}).$$

On the other hand,

$$\mathsf{SC}_{\lambda^d}(\widetilde{n}, 2, \mathbf{P}) = \mathsf{SC}_{(\lambda+u)^d}(\widetilde{n}, 2, \widetilde{\mathbf{P}}),$$

where $\widetilde{\mathbf{P}}$ is a mixed strategy profile associated with the instance $\langle \widetilde{n}, 2 \rangle$ that assigns v pure players to link 2 and r mixed players to both links. Since the function $(\lambda + u)^d$ is a linear combination of monomials in λ with non-negative coefficients, we are reduced to Case (a). Hence, it follows that $\mathsf{SC}_{(\lambda+u)^d}(\widetilde{n}, 2, \widetilde{\mathbf{P}}) \leq \mathsf{SC}_{(\lambda+u)^d}(\widetilde{n}, 2, \widetilde{\mathbf{F}})$. Hence, this implies that $\mathsf{SC}_{\lambda^d}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\lambda^d}(n, 2, \mathbf{Q})$, as needed. The proof for the case $u > 0$ that $\mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{P}) \leq \mathsf{SC}_{\pi_d(\lambda)}(n, 2, \mathbf{F})$ is now complete.

Since we examined all possible cases, the proof is now complete. ∎

We now turn to the case of m identical links. We prove that the polynomial social cost of any Nash equilibrium is upper bounded by $(1 + \frac{1}{n-1})^d$ times the polynomial social cost of the fully mixed Nash equilibrium.

Theorem 4.44. *Consider the model of identical players and identical links. Fix an instance $\langle n, m \rangle$ with associated Nash equilibrium $\mathbf{P}$ and fully mixed Nash equilibrium $\mathbf{F}$. Then,*

$$\mathsf{SC}_{\pi_d(\lambda)}(n, m, \mathbf{P}) \leq \left(1 + \frac{1}{n-1}\right)^d \cdot \mathsf{SC}_{\pi_d(\lambda)}(n, m, \mathbf{F}) \ .$$

Proof. We first consider the case of the monomial latency cost function $\pi_d(\lambda) = \lambda^d$. We will later reduce the general case to this case.

Denote $\alpha = \frac{n}{m}$. For each link $j \in [m]$, denote $r_j = |\{i \in [n] : p_{ij} > 0\}|$. Assume, without loss of generality, that for each link $j \in [m]$, $r_j \geq 1$. Clearly, the average probability on link j is $\frac{\Lambda_j(\mathbf{P})}{r_j}$. Define $\beta_j = |\Lambda_j(\mathbf{P}) - \alpha|$. Roughly speaking, β_j is the excess expected latency on link j from the fair share α. Partition the set of links $[m]$ into

$$\mathcal{M}_1 = \{j \in [m] \mid 0 < \Lambda_j(\mathbf{P}) \leq \alpha\} \ ,$$
$$\mathcal{M}_2 = \{j \in [m] \mid \Lambda_j(\mathbf{P}) > \alpha\} \ .$$

Clearly, $\Lambda_j(\mathbf{P}) = \alpha - \beta_j$ for $j \in \mathcal{M}_1$ and $\Lambda_j(\mathbf{P}) = \alpha + \beta_j$ for $j \in \mathcal{M}_2$. Define now $q_j = \min_{i \in [n]} \{p_{ij} \mid p_{ij} > 0\}$. Clearly, $q_j \leq \frac{\Lambda_j(\mathbf{P})}{r_j}$.

Define $\beta = \max_{j \in \mathcal{M}_1} \beta_j$. We prove a simple fact.

Lemma 4.45. *For each link $j \in \mathcal{M}_2$ with $r_j \geq 2$, $\beta_j \leq \frac{\alpha - r_j \beta}{r_j - 1}$.*

Proof. Fix a link $j \in \mathcal{M}_2$ with $r_j \geq 2$ and a player $i_0 \in \{i \in [n] \mid p_{ij} > 0\}$ such that $p_{i_0 j} = q_j$. Consider a link $j' \in \mathcal{M}_1$ such that $\beta_{j'} = \beta$. Since $\mathbf{P}$ is a Nash equilibrium, $\lambda_{i_0 j}(\mathbf{P}) \leq \lambda_{i_0 j'}(\mathbf{P})$. Clearly, $\lambda_{i_0 j}(\mathbf{P}) = \Lambda_j(\mathbf{P}) - q_j + 1 = \alpha + \beta_j - q_j + 1$. Also, $\lambda_{i_0 j'}(\mathbf{P}) = \Lambda_{j'}(\mathbf{P}) - p_{i_0 j'} + 1 \leq \Lambda_{j'}(\mathbf{P}) + 1 = \alpha - \beta + 1$. It follows that $\alpha + \beta_j - q_j + 1 \leq \alpha - \beta + 1$ or $\beta_j + \beta \leq q_j \leq \frac{\Lambda_j(\mathbf{P})}{r_j} = \frac{\alpha + \beta_j}{r_j}$. This implies that $\beta_j \leq \frac{\alpha - r_j \beta}{r_j - 1}$, as needed. ∎

On one hand, Lemma 2.3 and Proposition 2.2 imply that

$$\begin{aligned}
&\mathsf{SC}_{\lambda^d}(n, m, \mathbf{P}) \\
&= \sum_{j \in [m]} \mathsf{BF}\left(\langle p_{1j}, \ldots, p_{nj}\rangle, \lambda^d\right) \\
&= \sum_{j \in \mathcal{M}_1} \mathsf{BF}\left(\langle p_{1j}, \ldots, p_{nj}\rangle, \lambda^d\right) + \sum_{j \in \mathcal{M}_2} \mathsf{BF}\left(\langle p_{1j}, \ldots, p_{nj}\rangle, \lambda^d\right) \\
&\leq \sum_{j \in \mathcal{M}_1} \mathsf{BF}\left(\frac{\alpha - \beta_j}{r_j}, r_j, \lambda^d\right) + \sum_{j \in \mathcal{M}_2} \mathsf{BF}\left(\frac{\alpha + \beta_j}{r_j}, r_j \lambda^d\right) \\
&= \sum_{j \in \mathcal{M}_1} \sum_{t \in [d]} \left(\frac{\alpha - \beta_j}{r_j}\right)^t \cdot S(d,t) \cdot (r_j)^{\underline{t}} + \sum_{j \in \mathcal{M}_2} \sum_{t \in [d]} \left(\frac{\alpha + \beta_j}{r_j}\right)^t \cdot S(d,t) \cdot (r_j)^{\underline{t}} \\
&= \sum_{t \in [d]} S(d,t) \cdot \left(\sum_{j \in \mathcal{M}_1} \left(\frac{\alpha - \beta_j}{r_j}\right)^t \cdot (r_j)^{\underline{t}} + \sum_{j \in \mathcal{M}_2} \left(\frac{\alpha + \beta_j}{r_j}\right)^t \cdot (r_j)^{\underline{t}}\right) .
\end{aligned}$$

On the other hand Lemma 4.41 and Proposition 2.2 imply that

$$\begin{aligned}
\mathsf{SC}_{\lambda^d}(n, m, \mathbf{F}) &= \sum_{j \in [m]} \mathsf{BF}\left(\langle f_{1j}, \ldots, f_{nj}\rangle, \lambda^d\right) \\
&= m \cdot \mathsf{BF}\left(\frac{1}{m}, n, \lambda^d\right) \\
&= \sum_{t \in [d]} S(d,t) \cdot \left(m \cdot \alpha^t \cdot \frac{n^{\underline{t}}}{n^t}\right) .
\end{aligned}$$

For each integer $t \in [d]$ define the function

$$\begin{aligned}
\Delta(t) = &\left(\frac{n}{n-1}\right)^t \left(m \cdot \alpha^t \cdot \frac{n^{\underline{t}}}{n^t}\right) \\
&- \left(\sum_{j \in \mathcal{M}_1} \left(\frac{\alpha - \beta_j}{r_j}\right)^t \cdot (r_j)^{\underline{t}} + \sum_{j \in \mathcal{M}_2} \left(\frac{\alpha + \beta_j}{r_j}\right)^t \cdot (r_j)^{\underline{t}}\right) .
\end{aligned}$$

We prove:

Lemma 4.46. *For each integer* $t \geq 1$, $\Delta(t) \geq 0$.

Proof. By induction on t. For the basis case, let $t = 1$. Then,

$$\begin{aligned}
\Delta(t) &= m \cdot \alpha - \left(\sum_{j \in \mathcal{M}_1} (\alpha - \beta_j) + \sum_{j \in \mathcal{M}_2} (\alpha + \beta_j)\right) \\
&= m \cdot \alpha - \left(\sum_{j \in \mathcal{M}_1} \Lambda_j(\mathbf{P}) + \sum_{j \in \mathcal{M}_2} \Lambda_j(\mathbf{P})\right) \\
&= m \cdot \alpha - n \\
&= 0 ,
\end{aligned}$$

as needed.

Assume inductively that the claim holds for $(t-1)$, for some integer $t \geq 2$. For the induction step, we will prove the claim for t.

Since $(r_j)^{\underline{t}} = 0$ for each $j \in \mathcal{M}_1 \cup \mathcal{M}_2$ with $r_j < t$, it follows that

$$\begin{aligned}
&\sum_{j\in\mathcal{M}_1} \left(\frac{\alpha-\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} + \sum_{j\in\mathcal{M}_2} \left(\frac{\alpha+\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} \\
&= \sum_{j\in\mathcal{M}_1: r_j\geq t} \left(\frac{\alpha-\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} + \sum_{j\in\mathcal{M}_2: r_j\geq t} \left(\frac{\alpha+\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} \\
&= \sum_{j\in\mathcal{M}_1: r_j\geq t} (\alpha-\beta_j)\frac{r_j-(t-1)}{r_j} \left(\frac{\alpha-\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j\in\mathcal{M}_2: r_j\geq t} (\alpha+\beta_j)\frac{r_j-(t-1)}{r_j} \left(\frac{\alpha+\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}}.
\end{aligned}$$

Note that for each link $j \in [m]$, the fraction $\frac{r_j-(t-1)}{r_j}$ is monotonically increasing in r_j (since $t \geq 2$). Since for each $j \in [m]$, $r_j \leq n$ and $\beta_j \geq 0$, it follows that

$$\begin{aligned}
&\sum_{j\in\mathcal{M}_1} \left(\frac{\alpha-\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} + \sum_{j\in\mathcal{M}_2} \left(\frac{\alpha+\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} \\
&\leq \sum_{j\in\mathcal{M}_1: r_j\geq t} \alpha\frac{n-(t-1)}{n} \left(\frac{\alpha-\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j\in\mathcal{M}_2: r_j\geq t} (\alpha+\beta_j)\frac{r_j-(t-1)}{r_j} \left(\frac{\alpha+\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}}. \qquad (4.21)
\end{aligned}$$

We proceed by case analysis.

1. Assume that for each link $j \in \mathcal{M}_2$ with $r_j \geq t$, $\beta_j \leq \frac{n-r_j}{m(r_j-1)}$. The implies that for each link $j \in \mathcal{M}_2$ with $r_j \geq t$, $\alpha + \beta_j \leq \frac{(n-1)r_j}{m(r_j-1)}$. Then, with (4.21),

$$\begin{aligned}
&\sum_{j\in\mathcal{M}_1} \left(\frac{\alpha-\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} + \sum_{j\in\mathcal{M}_2} \left(\frac{\alpha+\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} \\
&\leq \sum_{j\in\mathcal{M}_1: r_j\geq t} \alpha\frac{n-(t-1)}{n} \left(\frac{\alpha-\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j\in\mathcal{M}_2: r_j\geq t} \frac{(n-1)r_j}{m(r_j-1)}\frac{r_j-(t-1)}{r_j} \left(\frac{\alpha+\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}}
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{j \in \mathcal{M}_1 : r_j \geq t} \alpha \frac{n-(t-1)}{n} \left(\frac{\alpha - \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j \in \mathcal{M}_2 : r_j \geq t} \frac{n-1}{m} \frac{r_j-(t-1)}{r_j - 1} \left(\frac{\alpha + \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&\leq \sum_{j \in \mathcal{M}_1 : r_j \geq t} \alpha \frac{n-(t-1)}{n} \left(\frac{\alpha - \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j \in \mathcal{M}_2 : r_j \geq t} \frac{n-1}{m} \frac{n-(t-1)}{n - 1} \left(\frac{\alpha + \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&= \sum_{j \in \mathcal{M}_1 : r_j \geq t} \alpha \frac{n-(t-1)}{n} \left(\frac{\alpha - \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j \in \mathcal{M}_2 : r_j \geq t} \alpha \frac{n-(t-1)}{n} \left(\frac{\alpha + \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \\
&= \alpha \frac{n-(t-1)}{n} \cdot \\
&\qquad \cdot \left(\sum_{j \in \mathcal{M}_1 : r_j \geq t} \left(\frac{\alpha - \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} + \sum_{j \in \mathcal{M}_2 : r_j \geq t} \left(\frac{\alpha + \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \right) \\
&= \alpha \frac{n-(t-1)}{n} \left(\sum_{j \in \mathcal{M}_1} \left(\frac{\alpha - \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} + \sum_{j \in \mathcal{M}_2} \left(\frac{\alpha + \beta_j}{r_j} \right)^{t-1} (r_j)^{\underline{t-1}} \right) \\
&\leq \alpha \frac{n-(t-1)}{n} \left(\frac{n}{n-1} \right)^{t-1} \left(m \cdot \alpha^{t-1} \cdot \frac{n^{\underline{t-1}}}{n^{t-1}} \right) \\
&< \left(\frac{n}{n-1} \right)^{t} \left(m \alpha^{t} \frac{n^{\underline{t}}}{n^{t}} \right) ,
\end{aligned}
$$

where the induction hypothesis was used for the last inequality. This implies that $\Delta(t) \geq 0$, as needed.

2. Assume now that there exists a link $j \in \mathcal{M}_2$ with $r_j \geq t$ such that $\beta_j > \frac{n-r_j}{m(r_j-1)}$. By Lemma 4.45, for each link $j \in \mathcal{M}_2$ with $r_j \geq t \geq 2$,

$$
\begin{aligned}
(\alpha + \beta_j) \cdot \frac{r_j - (t-1)}{r_j} &\leq \left(\alpha + \frac{\alpha - r_j \beta}{r_j - 1} \right) \cdot \frac{r_j - (t-1)}{r_j} \\
&\overset{\beta \geq 0}{\leq} \frac{\alpha r_j}{r_j - 1} \cdot \frac{r_j - (t-1)}{r_j} \\
&= \alpha \cdot \frac{r_j - (t-1)}{r_j - 1} \\
&\leq \alpha \cdot \frac{n-(t-1)}{n-1} .
\end{aligned}
$$

Thus, with (4.21),

$$\begin{aligned}
&\sum_{j\in\mathcal{M}_1}\left(\frac{\alpha-\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} + \sum_{j\in\mathcal{M}_2}\left(\frac{\alpha+\beta_j}{r_j}\right)^t (r_j)^{\underline{t}} \\
&\leq \sum_{j\in\mathcal{M}_1 : r_j \geq t} \alpha\frac{n-(t-1)}{n}\left(\frac{\alpha-\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} \\
&\quad + \sum_{j\in\mathcal{M}_2 : r_j \geq t} \alpha\frac{n-(t-1)}{n-1}\left(\frac{\alpha+\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} \\
&\leq \frac{n}{n-1}\alpha\frac{n-(t-1)}{n} \cdot \\
&\qquad \cdot\left(\sum_{j\in\mathcal{M}_1 : r_j \geq t}\left(\frac{\alpha-\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}} + \sum_{j\in\mathcal{M}_2 : r_j \geq t}\left(\frac{\alpha+\beta_j}{r_j}\right)^{t-1} (r_j)^{\underline{t-1}}\right) \\
&\leq \frac{n}{n-1}\alpha\frac{n-(t-1)}{n}\left(\frac{n}{n-1}\right)^{t-1}\left(m\cdot\alpha^{t-1}\frac{n^{\underline{t-1}}}{n^{t-1}}\right) \\
&= \left(\frac{n}{n-1}\right)^t\left(m\alpha^t\frac{n^{\underline{t}}}{n^t}\right),
\end{aligned}$$

where the induction hypothesis was used for the last inequality. This implies that $\Delta(t) \geq 0$, as needed.

This completes the proof of the Lemma 4.46. ■

Lemma 4.46 implies that $\mathsf{SC}_{\lambda^d}(n,m,\mathbf{P}) \leq \left(1+\frac{1}{n-1}\right)^d \cdot \mathsf{SC}_{\lambda^d}(n,m,\mathbf{F})$. Hence,

$$\begin{aligned}
\mathsf{SC}_{\pi_d(\lambda)}(n,m,\mathbf{P}) &= \sum_{0\leq t\leq d} a_t \cdot \mathsf{SC}_{\lambda^t}(n,m,\mathbf{P}) \\
&\leq \sum_{0\leq t\leq d} a_t \cdot \left(1+\frac{1}{n-1}\right)^t \cdot \mathsf{SC}_{\lambda^t}(n,m,\mathbf{F}) \\
&\leq \left(1+\frac{1}{n-1}\right)^d \cdot \sum_{0\leq t\leq d} a_t \cdot \mathsf{SC}_{\lambda^t}(n,m,\mathbf{F}) \\
&= \left(1+\frac{1}{n-1}\right)^d \cdot \mathsf{SC}_{\pi_d(\lambda)}(n,m,\mathbf{F}),
\end{aligned}$$

as needed. The proof of Theorem 4.44 is now complete. ■

We remark that the proof of Theorem 4.44 follows the proof of Theorem 4.42. However, it is more complicated in dealing with an arbitrary number of links.

In the process of publishing the corresponding paper in a journal, an unknown referee contributed a substantial improvement to Theorem 4.44, showing that the factor $\left(1+\frac{1}{n-1}\right)^d$ is *not* necessary.

Theorem 4.47 (Gairing et al. [43]). *Consider the model of identical players and identical links. Fix an instance* $\langle n, m\rangle$ *with associated Nash equilibrium* $\mathbf{P}$ *and fully mixed Nash equilibrium* $\mathbf{F}$*. Then,*

$$\mathsf{SC}_{\pi_d(\lambda)}(n, m, \mathbf{P}) \leq \mathsf{SC}_{\pi_d(\lambda)}(n, m, \mathbf{F})\,.$$

4.5.1.2 The Monomial and Polynomial Prices of Anarchy.

We are now ready to prove our upper bounds on the price of anarchy for monomial and polynomial social cost.

Identical Players and Identical Links

We first consider the model of identical players and identical links. We use Theorem 4.47 to prove upper bounds on the price of anarchy for monomial and polynomial social cost (Theorem 4.48 and Corollary 4.49).

Theorem 4.48. *Consider the model of identical players and identical links. Then,*

$$\mathsf{PoA}_{\lambda^d} \leq B_d\,.$$

Proof. Fix any instance $\langle n, m\rangle$ with an associated fully mixed Nash equilibrium $\mathbf{F}$. By Proposition 2.2,

$$\begin{aligned}
\mathsf{SC}_{\lambda^d}(n, m, \mathbf{F}) &= \sum_{j\in[m]} \mathsf{BF}\left(\langle f_{1j}, \ldots, f_{nj}\rangle, \lambda^d\right)\\
&= m \cdot \mathsf{BF}\left(\langle f_{1j}, \ldots, f_{nj}\rangle, \lambda^d\right)\\
&= m \cdot \sum_{t\in[d]} \left(\frac{1}{m}\right)^t \cdot S(d,t) \cdot n^{\underline{t}}\\
&\leq m \cdot \sum_{t\in[d]} \left(\frac{1}{m}\right)^t \cdot S(d,t) \cdot n^{t}\,.
\end{aligned}$$

We now proceed by case analysis.

1. Assume first that $n \geq m$. Recall that in this case, $\mathsf{OPT}_{\lambda^d}(\mathbf{w}, m) \geq m \cdot \left(\frac{n}{m}\right)^d$. Hence,

$$\begin{aligned}
\frac{\mathsf{SC}_{\lambda^d}(n, m, \mathbf{F})}{\mathsf{OPT}_{\lambda^d}(n, m)} &\leq \frac{1}{m} \cdot \left(\frac{m}{n}\right)^d \cdot m \cdot \sum_{t\in[d]} \left(\frac{1}{m}\right)^t \cdot S(d,t) \cdot n^t\\
&= \sum_{t\in[d]} \left(\frac{m}{n}\right)^{d-t} \cdot S(d,t)\\
&\leq \sum_{t\in[d]} S(d,t)\\
&= B_d\,.
\end{aligned}$$

2. Assume now that $n < m$. Recall that, in this case, $\mathsf{OPT}_{\lambda^d}(n, m) \geq n$. Hence,

$$\begin{aligned}\frac{\mathsf{SC}_{\lambda^d}(n, m, \mathbf{F})}{\mathsf{OPT}_{\lambda^d}(n, m)} &\leq \frac{1}{n} m \cdot \sum_{t \in [d]} \left(\frac{1}{m}\right)^t \cdot S(d, t) \cdot n^t \\ &= \sum_{t \in [d]} \left(\frac{n}{m}\right)^{t-1} \cdot S(d, t) \\ &\leq \sum_{t \in [d]} S(d, t) \\ &= B_d \,.\end{aligned}$$

So, in all cases, $\frac{\mathsf{SC}_{\lambda^d}(n,m,\mathbf{F})}{\mathsf{OPT}_{\lambda^d}(n,m)} \leq B_d$. Theorem 4.47 implies now the claim. ■

Note that the upper bound on the monomial price of anarchy established in Theorem 4.48 approaches B_d as n approaches infinity. By Lemma 4.40 and since the Bell numbers are increasing in their order, Theorem 4.48 immediately implies:

Corollary 4.49. *Consider the model of identical players and identical links. Then,*

$$\mathsf{PoA}_{\pi_d(\lambda)} \leq B_d \,.$$

Identical Players and Two Identical Links

We now turn to the model of identical players and *two* identical links. Again, we use Theorem 4.47 to prove (tight) upper bounds on the price of anarchy for monomial and polynomial social cost (Theorem 4.50 and Corollary 4.51).

Theorem 4.50. *Consider the model of identical players and two identical links. Then,*

$$\mathsf{PoA}_{\lambda^d} \leq 2^{d-2} \left(1 + \left(\frac{1}{n}\right)^{d-1}\right) .$$

This bound is tight for $n = 2$.

Proof. We start with the upper bound. Fix any instance $\langle n, 2 \rangle$ with an associated Nash equilibrium $\mathbf{P}$. On the one hand, by Theorem 4.47,

$$\begin{aligned}
\mathsf{SC}_{\lambda^d}(n,2,\mathbf{P}) &\le \mathsf{SC}_{\lambda^d}(n,2,\mathbf{F}) \\
&= 2\cdot \mathsf{BF}(\frac{1}{2},n,\lambda^d) \\
&= 2\cdot \sum_{t\in[d]} \left(\frac{1}{2}\right)^t \cdot S(d,t)\cdot n^{\underline{t}} \\
&\le 2\cdot \left(\frac{1}{2}\cdot S(d,1)\cdot n + \frac{1}{4}\cdot \sum_{2\le i\le d} S(d,i) n^{\underline{i}}\right) \\
&= 2\cdot \left(\frac{1}{2}\cdot S(d,1)\cdot n + \frac{1}{4}\cdot (n^d - S(d,1)\cdot n)\right) \\
&= 2\cdot \left(\frac{n}{4}+\frac{n^d}{4}\right) .
\end{aligned}$$

On the other hand,

$$\mathsf{OPT}_{\lambda^d}(n,2) \ge 2\cdot \left(\frac{n}{2}\right)^d .$$

It follows that

$$\begin{aligned}
\frac{\mathsf{SC}_{\lambda^d}(n,2,\mathbf{F})}{\mathsf{OPT}_{\lambda^d}(n,2)} &\le \left(\frac{2}{n}\right)^d \cdot \left(\frac{n}{4}+\frac{n^d}{4}\right) \\
&= 2^{d-2}\left(1+\left(\frac{1}{n}\right)^{d-1}\right) ,
\end{aligned}$$

as needed. To prove that the upper bound is tight for $n = 2$, note that it becomes $2^{d-2} + \frac{1}{2}$. We continue to prove that this is also a lower bound for $n = 2$. Fix an instance $\langle 2, 2\rangle$. Then,

$$\mathsf{OPT}_{\lambda^d}(2,2) = 2 ,$$

while

$$\begin{aligned}
\mathsf{SC}_{\lambda^d}(2,2,\mathbf{F}) &= 2\cdot \sum_{t\in[d]} \left(\frac{1}{2}\right)^t \cdot S(d,t)\cdot 2^{\underline{t}} \\
&= 2\cdot \left(\frac{1}{2}\cdot S(d,1)\cdot 2 + \frac{1}{4}\cdot S(d,2)\cdot 2\cdot 1\right) \\
&= 2\cdot \left(S(d,1) + \frac{1}{2}\cdot S(d,2)\right) \\
&= 2\cdot \left(1 + \frac{1}{2}\cdot (2^{d-1}-1)\right) \\
&= 2\cdot \left(2^{d-2}+\frac{1}{2}\right) .
\end{aligned}$$

It follows that $\mathsf{PoA}_{\lambda^d} \ge 2^{d-2} + \frac{1}{2}$ as needed. ■

By Lemma 4.40, Theorem 4.50 immediately implies:

Corollary 4.51. *Consider the model of identical players and two identical links. Then,*

$$\mathsf{PoA}_{\pi_d(\lambda)} \leq 2^{d-2}\left(1+\left(\frac{1}{n}\right)^{d-1}\right).$$

This bound is tight for $n = 2$.

Proof. We will show that for all integer $t \in [2, d]$ the upper bound on PoA_{λ^t} is larger than the upper bound on $\mathsf{PoA}_{\lambda^{t-1}}$. Clearly,

$$\begin{aligned}
&2^{t-2}\left(1+\left(\frac{1}{n}\right)^{t-1}\right) - 2^{t-3}\left(1+\left(\frac{1}{n}\right)^{t-2}\right)\\
&= 2^{t-3}\left(2+\frac{2}{n}\left(\frac{1}{n}\right)^{t-2}\right) - 2^{t-3}\left(1+\left(\frac{1}{n}\right)^{t-2}\right)\\
&= 2^{t-3}\left(1-\left(1-\frac{2}{n}\right)\left(\frac{1}{n}\right)^{t-2}\right)\\
&\geq 2^{t-3}\left(1-\left(1-\frac{2}{n}\right)\right)\\
&= \frac{2^{t-2}}{n}\\
&> 0\,,
\end{aligned}$$

as needed. Tightness, for $n = 2$, follows from the tightness of Theorem 4.50 ∎

4.6 Conclusion and Discussion

In this chapter, we have studied routing games on parallel links. For this setting, we have provided many results concerning the computational complexity of pure Nash equilibria. Moreover, we proved an extensive collection of results related to the price of anarchy in various sub-models. Although, routing games on parallel links have received a lot of attention, many problems are still tantalizing open. We only state some of them.

- For the model of identical links, our nashification algorithm NashifyIdentical is based only on selfish steps. Is it also possible to provide a polynomial-time nashification algorithm, for the model of *related* links, that solely depends on selfish steps?
- We have described a polynomial-time algorithm to compute a pure Nash equilibrium for the model of restricted strategy sets and identical links. Is it possible to provide such an algorithm for the model of restricted strategy sets and *related* links?
- Our bounds on the price of anarchy for polynomial social cost are all for the model of identical players. Proving such bounds for *arbitrary* players remains a challenging open problem.

5

Weighted Congestion Games

5.1 Introduction

In this chapter, we present strong results on the price of anarchy for *congestion games* and *weighted congestion games*. Such games have been formally introduced in Section 3.2. In a congestion game, there is a set of *resources* and the strategy set of each player is a subset of the power set of the resources. Thus, a pure strategy might consist of multiple resources. This stands in contrast to the games studied in Chapter 4 where each pure strategy consists of a single resource (link). For each resource, there is a *latency function* which describes the latency of this resource. In this chapter, we allow for polynomial latency functions with maximum degree d and non-negative coefficients. Each player aims to minimize its *private cost* which is defined as the (expected) sum of the latencies of its chosen resources. For (unweighted) congestion games the latency of a resource only depends on the number of players sharing this resource. In a *weighted* congestion game, players have weights and thus different influence on the congestion of the resources. Weighted congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem. A typical resource sharing problem is that of routing. In a routing game the strategy sets of the players correspond to paths in a network. Routing games where the demand of the players cannot be split among multiple paths are also called *(weighted) network congestion games*.

5.1.1 Summary of Results

In this chapter, we prove *exact* bounds on the price of anarchy for unweighted and weighted congestion games with polynomial latency functions. We use the total latency as social cost measure. This improves on results by Awerbuch et al. [5] and Christodoulou and Koutsoupias [17], where non-matching upper and lower bounds are given.
We now describe our findings in more detail.

- For *unweighted congestion games* we show that the price of anarchy ($\mathsf{PoA}_{\mathsf{TL}}$) is exactly

$$\mathsf{PoA_{TL}} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$ and Φ_d is a natural generalization of the golden ratio to larger dimensions such that Φ_d is the solution to $(\Phi_d + 1)^d = \Phi_d^{d+1}$. Prior to this paper the best known upper and lower bounds were shown to be of the form $d^{d(1-o(1))}$ [17]. However, the term $o(1)$ still hides a gap between the upper and the lower bound.

- For *weighted congestion games* we show that the price of anarchy ($\mathsf{PoA_{TL}}$) is exactly

$$\mathsf{PoA_{TL}} = \Phi_d^{d+1}.$$

This result closes the gap between the so far best upper and lower bounds of $O(2^d d^{d+1})$ and $\Omega(d^{d/2})$ from [5].

We show that the above values on the price of anarchy also hold for the subclasses of unweighted and weighted *network* congestion games.

For our upper bounds we use a similar analysis as in [17]. The core of our analysis is to determine parameters c_1 and c_2 such that

$$y \cdot f(x+1) \le c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y) \tag{5.1}$$

for all polynomial latency functions f of maximum degree d and for all reals $x, y \ge 0$. For the case of unweighted demands it suffices to show (5.1) for all integers x, y. In order to prove their upper bound Christodoulou and Koutsoupias [17] looked at (5.1) with $c_1 = \frac{1}{2}$ and gave an asymptotic estimate for c_2. In our analysis we optimize both parameters c_1, c_2. This optimization process requires new ideas and is non-trivial.

		unweighted $\mathsf{PoA_{TL}}$			weighted $\mathsf{PoA_{TL}}$	
d	Φ_d	Our exact result	Upper Bound [17]	Lower bound [17]	Our exact result	Lower bound [5]
1	1.618	**2.5**	2.5	2.5	**2.618**	2.618
2	2.148	**9.583**	10	(2.5)	**9.909**	(2.618)
3	2.630	**41.54**	47	(2.5)	**47.82**	5
4	3.080	**267.6**	269	21.33	**277.0**	15
5	3.506	**1,514**	2,154	42.67	**1,858**	52
6	3.915	**12,345**	15,187	85.33	**14,099**	203
7	4.309	**98,734**	169,247	170.7	**118,926**	877
8	4.692	**802,603**	1,451,906	14,762	**1,101,126**	4,140
9	5.064	**10,540,286**	20,241,038	44,287	**11,079,429**	21,147
10	5.427	**88,562,706**	202,153,442	132,860	**120,180,803**	115,975

Table 5.1. Comparison of our results to [17] and [5]

Table 5.1 shows a numerical comparison of our bounds with the previous results of Awerbuch et al. [5] and Christodoulou and Koutsoupias [17].

For $d \ge 2$, the table only gives the respective lower bounds that are given in the cited works (before any estimates are applied). Values in parentheses denote cases in which the bound for linear functions is better than the general case.

In [5, Theorem 4.3], a construction scheme for networks is described with price of anarchy approximating $\frac{1}{e}\sum_{k=1}^{\infty}\frac{k^d}{k!}$ which yields the d-th Bell number. In [17, Theorem 10], a network with price of anarchy $\frac{(N-1)^{d+2}}{N}$ is given, with N being the largest integer for which $(N-1)^{d+2} \leq N^d$ holds.

The column with the upper bound from [17] is computed by using (5.1) with $c_1 = \frac{1}{2}$ and optimizing c_2 with help of our analysis. Thus, the column shows the best possible bounds that can be shown with $c_1 = \frac{1}{2}$.

5.1.2 Related Work

The papers most closely related to our work are those of Awerbuch et al. [5] and Christodoulou and Koutsoupias [17, 16]. For (unweighted) congestion games and social cost defined as average private cost (which in this case is the same as total latency) it has been shown that the price of anarchy of pure Nash equilibria is $\frac{5}{2}$ for linear latency functions and $d^{\Theta(d)}$ for polynomial latency functions of maximum degree d [5, 17]. The bound of $\frac{5}{2}$ for linear latency function also holds for the (mixed) price of anarchy [16]. For *weighted* congestion games and social cost defined as the total latency, the (mixed) price of anarchy is $\frac{3+\sqrt{5}}{2}$ in case of linear latency functions and $d^{\Theta(d)}$ in case of polynomial latency functions [5].

Since the routing games on parallel links are a special class of weighted congestion games, all results concerning the price of anarchy that are either described in Section 4.1.2 or presented throughout Chapter 4 belong to the related work here. Of particular interest is the paper of Lücking et al. [71], where they studied the total latency (they call it quadratic social cost) for routing games on parallel links. For this model, Lücking et al. [71] showed that the price of anarchy is exactly $\frac{4}{3}$ for the case of identical players and related links and $\frac{9}{8}$ for the case of arbitrary players and identical links.

The class of *congestion games* has been introduced by Rosenthal [88] and extensively studied afterwards (see e.g. [1, 2, 27, 38, 39, 77, 78, 96]). In Rosenthal's model the strategy of each player is a subset of resources. Resource utility functions can be arbitrary but they only depend on the number of players sharing the same resource. Rosenthal showed that such games always admit a pure Nash equilibrium using a potential function. Monderer and Shapley [78] and later Voorneveld et al. [96] characterized games that possess a potential function as potential games and showed their relation to congestion games. Milchtaich [77] considers weighted congestion games with player-specific payoff functions and shows that these games do not admit a pure Nash equilibrium in general. Ackermann et al. [2] further studied the existence of pure Nash equilibria in weighted congestion games and (unweighted) congestion games with player-specific payoff functions. Fotakis et al. [38, 39] considered the price of anarchy for symmetric weighted network congestion games in layered networks [38] and for symmetric (unweighted) network congestion games in general networks [39]. In both cases they defined social cost as expected maximum latency. The complexity of computing a pure Nash equilibrium has been studied by Fabrikant et al. [27]. On the one hand, they proved that for any symmetric network congestion game, a pure Nash equilibrium can be constructed in polynomial time, by computing the optimum of Rosenthal's potential function. This is accomplished through a nice reduction to an instance of the min-cost flow problem. On the other hand, Fabrikant et al. [27] showed that the problem of computing a pure Nash equilibrium becomes PLS-complete, if we allow for asymmetric or non-network congestion games. Ackermann et al. [1] further studied the complexity of computing a pure Nash equilibrium in congestion games. In particular, they showed that a pure Nash equilibrium can be computed in polynomial time, if the strategy sets of the players have a certain property. For a survey on weighted congestion games we refer to [48].

Inspired by the arisen interest in the price of anarchy Roughgarden and Tardos [92] re-investigated the Wardrop model and used the *total latency* as a social cost measure. In this context the price of anarchy was shown to be $\frac{4}{3}$ for linear latency functions [92] and $\Theta(\frac{d}{\log d})$ [89] for polynomial latency functions of maximum degree d. An overview on results for this model can be found in the recent book of Roughgarden [90].

5.1.3 Organization

The rest of this chapter is organized as follows. We present our exact bounds on the price of anarchy for unweighted congestion games in Section 5.2 and for weighted congestion games in Section 5.3. We conclude in Section 5.4 with a discussion on our results.

5.2 Price of Anarchy for Unweighted Congestion Games

In this section, we prove the exact value for the price of anarchy of *unweighted* congestion games with polynomial latency functions. We start by showing the upper bound in Section 5.2.1. In Section 5.2.2 we provide a matching lower bound.

5.2.1 Upper Bound

Before we can state our upper bound on the price of anarchy for unweighted congestion games, we introduce two technical lemmas (Lemma 5.1 and Lemma 5.2). These lemmas are crucial for determining c_1 and c_2 in (5.1) and thus for proving the upper bound in Theorem 5.3.

Lemma 5.1. *Let* $0 \le c < 1$ *and* $d \in \mathbb{N}_0$ *then*

$$\max_{x\in\mathbb{N}_0, y\in\mathbb{N}} \left\{ \left(\frac{x+1}{y}\right)^d - c \cdot \left(\frac{x}{y}\right)^{d+1} \right\} = \max_{x\in\mathbb{N}_0} \left\{ (x+1)^d - c \cdot x^{d+1} \right\}.$$

Proof. Let

$$g(x, y, c) = \left(\frac{x+1}{y}\right)^d - c \cdot \left(\frac{x}{y}\right)^{d+1}.$$

We will show that for all $x \in \mathbb{N}_0$, $y \in \mathbb{N}$ there exists $\hat{x} \in \mathbb{N}_0$ such that

$$g(\hat{x}, 1, c) \geq g(x, y, c) \quad \forall 0 \leq c < 1.$$

Let $x \in \mathbb{N}_0$, $y \in \mathbb{N}$ be arbitrary non-negative integers. If $y \geq x+1$ then we can choose $\hat{x} = 0$ to see that $g(0, 1, c) = 1 \geq g(x, y, c)$. So in the following we may assume that $y \leq x$.

Let $\hat{x}$ be the smallest integer such that $g(\hat{x}, 1, 0) \geq g(x, y, 0)$, that is

$$\hat{x} = \left\lceil \frac{x+1-y}{y} \right\rceil.$$

To complete the proof, we will show that also $g(\hat{x}, 1, 1) \geq g(x, y, 1)$, or equivalently

$$\begin{aligned}
(\hat{x}+1)^d - \hat{x}^{d+1} &\geq \left(\frac{x+1}{y}\right)^d - \left(\frac{x}{y}\right)^{d+1} \\
\Leftrightarrow \left\lceil \frac{x+1}{y} \right\rceil^d - \left\lceil \frac{x+1-y}{y} \right\rceil^{d+1} &\geq \left(\frac{x+1}{y}\right)^d - \left(\frac{x}{y}\right)^{d+1} \\
\Leftarrow \qquad \left(\frac{x}{y}\right)^{d+1} &\geq \left\lceil \frac{x+1-y}{y} \right\rceil^{d+1} \\
\Leftrightarrow \qquad \frac{x}{y} &\geq \left\lceil \frac{x+1-y}{y} \right\rceil
\end{aligned}$$

To see that the last inequality holds, recall that $x \geq y$. Thus we can express x as $x = b_1 \cdot y + b_2$ where b_1 and b_2 are integers with $b_1 \geq 1$ and $0 \leq b_2 < y$. Then the last inequality reduces to $b_1 + \frac{b_2}{y} \geq b_1$, which is fulfilled. This completes the proof of the lemma. ■

Lemma 5.2. *Let* $d \in \mathbb{N}$ *and*

$$\mathcal{F}_d = \{g_r^{(d)} : \mathbb{R} \to \mathbb{R} \mid g_r^{(d)}(x) = (r+1)^d - x \cdot r^{d+1}, r \in \mathbb{R}_{\geq 0}\}$$

be an infinite set of linear functions. Furthermore, let $\gamma(s,t)$ *for* $s, t \in \mathbb{R}_{\geq 0}$ *and* $s \neq t$ *denote the intersection abscissa of* $g_s^{(d)}$ *and* $g_t^{(d)}$. *Then it holds for any* $s, t, u \in \mathbb{R}_{\geq 0}$ *with* $s < t < u$ *that* $\gamma(s,t) > \gamma(s,u)$ *and* $\gamma(u,s) > \gamma(u,t)$.

Proof. We first show that $\gamma(v, v+\delta)$ is strictly decreasing in $v \in \mathbb{R}_{\geq 0}$, for any $\delta \in \mathbb{R}_{>0}$. Afterwards we will show that this implies the lemma.

For some $v \in \mathbb{R}_{\geq 0}$, consider now the two linear functions $g_v^{(d)}$, $g_{v+\delta}^{(d)}$ from $\mathcal{F}_d$. They intersect at

$$\gamma(v, v+\delta) = \frac{(v+1+\delta)^d - (v+1)^d}{(v+\delta)^{d+1} - v^{d+1}}.$$

Computing the first derivative in v of $\gamma(v, v+\delta)$ yields

$$\begin{aligned}
\frac{\partial\gamma(v,v+\delta)}{\partial v} &= \frac{d\left[(v+1+\delta)^{d-1}-(v+1)^{d-1}\right]\cdot\left[(v+\delta)^{d+1}-v^{d+1}\right]}{((v+1)^{d+1}-v^{d+1})^2} \\
&\quad - \frac{(d+1)\left[(v+1+\delta)^{d}-(v+1)^{d}\right]\cdot\left[(v+\delta)^{d}-v^{d}\right]}{((v+1)^{d+1}-v^{d+1})^2} \\
&< \frac{d\left[(v+1+\delta)^{d-1}-(v+1)^{d-1}\right]\cdot\left[(v+\delta)^{d+1}-v^{d+1}\right]}{((v+1)^{d+1}-v^{d+1})^2} \\
&\quad - \frac{d\left[(v+1+\delta)^{d}-(v+1)^{d}\right]\cdot\left[(v+\delta)^{d}-v^{d}\right]}{((v+1)^{d+1}-v^{d+1})^2} \\
&= d\cdot\left[\frac{(1+\delta)(v+1+\delta)^{d-1}v^d+(1-\delta)(v+1)^{d-1}(v+\delta)^d}{((v+1)^{d+1}-v^{d+1})^2}\right. \\
&\quad \left. - \frac{(v+1+\delta)^{d-1}(v+\delta)^d+(v+1)^{d-1}v^d}{((v+1)^{d+1}-v^{d+1})^2}\right]
\end{aligned}$$

We now show (by induction over d) that

$$\begin{aligned}
&(1+\delta)(v+1+\delta)^{d-1}v^d+(1-\delta)(v+1)^{d-1}(v+\delta)^d \\
&\quad -(v+1+\delta)^{d-1}(v+\delta)^d-(v+1)^{d-1}v^d<0
\end{aligned} \tag{5.2}$$

and thus $\gamma(v, v+\delta)$ is strictly decreasing in v:

Clearly, (5.2) holds for $d = 1$ as

$$(1+\delta)v+(1-\delta)(v+\delta)-(v+\delta)-v=-\delta^2<0.$$

It also holds for $v = 0$ as

$$(1-\delta)\delta^d-(1+\delta)^{d-1}\delta^d<0.$$

Thus, we only consider $v > 0$ in the following. Assume that our induction hypothesis (5.2) holds for a natural d. We then multiply with $(v + 1 + \delta)v$ and get:

$$\begin{aligned}
&(1+\delta)(v+1+\delta)^d v^{d+1}+(1-\delta)(v+1)^{d-1}(v+\delta)^d(v+1+\delta)v \\
&\quad -\underbrace{(v+1+\delta)^d(v+\delta)^d}_{=A}v-\underbrace{(v+1)^{d-1}v^{d+1}}_{=B}(v+1+\delta)<0 \\
\Leftrightarrow\quad &(1+\delta)(v+1+\delta)^d v^{d+1}+\underbrace{(1-\delta)(v+1)^{d-1}(v+\delta)^d(v+1+\delta)v}_{=C} \\
&\quad -\delta B+\delta A-(v+1+\delta)^d(v+\delta)^{d+1}-(v+1)^d v^{d+1}<0.
\end{aligned}$$

Thus, if we define $D = (1 - \delta)(v + 1)^d (v + \delta)^{d+1}$, proving the inductive step $d \to d + 1$ reduces to showing that

$$C-\delta B+\delta A\geq D$$

or equivalently

$$C - \delta B + \delta A - D \geq 0.$$

Now,

$$\begin{aligned}
&C - \delta B + \delta A - D \\
&= (1-\delta)(v+1)^{d-1}(v+\delta)^d(v+1+\delta)v - \delta B + \delta A - D \\
&= (1-\delta)(v+1)^{d-1}(v+\delta)^d(v+1)v - \delta B + \delta A \\
&\quad + \delta(1-\delta)(v+1)^{d-1}(v+\delta)^d v - D \\
&= (1-\delta)(v+1)^{d}(v+\delta)^d(v+\delta) - \delta B + \delta A \\
&\quad + \delta(1-\delta)(v+1)^{d-1}(v+\delta)^d v - \delta(1-\delta)(v+1)^d(v+\delta)^d - D \\
&= D - \delta B + \delta A - \delta(1-\delta)(v+1)^{d-1}(v+\delta)^d - D \\
&= \delta\left(-B + A - (1-\delta)(v+1)^{d-1}(v+\delta)^d\right) \\
&> \delta\left(-B + A - (v+1)^{d-1}(v+\delta)^d\right) \\
&> \delta\left(-(v+1)^{d-1}(v+\delta)^d v + A - (v+1)^{d-1}(v+\delta)^d\right) \\
&= \delta\left(-(v+1)^{d}(v+\delta)^d + (v+1+\delta)^d(v+\delta)^d\right) \\
&> 0.
\end{aligned}$$

This last inequality obviously holds for any $\delta > 0$, thus $\gamma(v, v+\delta)$ is strictly decreasing in $v \in \mathbb{R}_{\geq 0}$, for any $\delta \in \mathbb{R}_{>0}$.

It follows that $\gamma(v + k \cdot \delta, v + (k+1) \cdot \delta)$ strictly decreases as $k \in \mathbb{Z}$, $k \geq -\frac{v}{\delta}$, becomes larger. We separately consider functions $g^{(d)}_{v+k\cdot\delta}$ for integers $k > 0$ and $k < 0$:

- We first turn to the case where $k > 0$. Then the intersection of $g^{(d)}_{v+k\cdot\delta}$ and $g^{(d)}_{v+(k+1)\cdot\delta}$ must lie above $g^{(d)}_v$. This can easily be seen due to the fact that $g^{(d)}_{v+k\cdot\delta}(x) > g^{(d)}_v(x)$ holds for any $x < \gamma(v, v+\delta)$ and any $k \in \mathbb{N}$. Now recall that the slope of $g^{(d)}_{v+k\cdot\delta}$ is $-(v + k \cdot \delta)^{d+1}$, which is decreasing as k becomes larger. Since the aforementioned intersection lies above $g^{(d)}_v$, we have that $\gamma(v, v + (k+1) \cdot \delta) < \gamma(v, v + k \cdot \delta)$.
- Now let $-\frac{v}{\delta} \leq k < 0$. For clarity set $j = -k$ and consider the intersection of $g^{(d)}_{v-j\cdot\delta}$ and $g^{(d)}_{v-(j+1)\cdot\delta}$ which still lies above $g^{(d)}_v$. Obviously, $g^{(d)}_{v-j\cdot\delta}(x) > g^{(d)}_v(x)$ holds for any $x > \gamma(v, v+\delta)$ and any $j \in \mathbb{N}$. With corresponding arguments as in the first case we get that $\gamma(v, v - j \cdot \delta) < \gamma(v, v - (j+1) \cdot \delta)$.

Hence, $\gamma(v, v + k \cdot \delta)$ is increasing in any $k \in \mathbb{Z} \setminus \{0\}$, $k \geq -\frac{v}{\delta}$, and the lemma holds for any $s < t < u$ where $(t-s)$ and $(u-s)$ are rational – by choosing δ as the reciprocal of a common denominator of $(t-s)$ and $(u-s)$. Note that this is also a denominator of $(u-t)$. Finally, as $x \mapsto \gamma(v, v+x)$ is continuous in any $x \in [-v, \infty) \setminus \{0\}$, the lemma follows. ■

We are now ready to prove our upper bound on the price of anarchy for unweighted congestion games with polynomial latency functions.

Theorem 5.3. *For unweighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have*

$$\mathsf{PoA}_{\mathsf{TL}} \leq \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}$$

where $k = \lfloor \Phi_d \rfloor$.

Proof. Let $\mathbf{P} = (P_1, ..., P_n)$ be a (mixed) Nash equilibrium and let $\mathbf{Q} = (Q_1, ..., Q_n)$ be a pure strategy profile with optimum social cost. Since $\mathbf{P}$ is a Nash equilibrium, player $i \in [n]$ cannot improve by switching from strategy P_i to strategy Q_i. Thus,

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{P}) &= \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in s_i} f_e(\delta_e(\mathsf{s})) \\
&\leq \mathsf{PC}_i(\mathbf{P}_{-i}, Q_i) \\
&= \sum_{\mathsf{s}' \in S_{-i}} p(\mathsf{s}') \left[\sum_{e \in Q_i \cap s_i} f_e(\delta_e(\mathsf{s})) + \sum_{e \in Q_i \setminus s_i} f_e(\delta_e(\mathsf{s}) + 1) \right] \\
&\leq \sum_{\mathsf{s}' \in S_{-i}} p(\mathsf{s}') \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + 1) \\
&= \sum_{s_i \in S_i} p(i, s_i) \sum_{\mathsf{s}' \in S_{-i}} p(\mathsf{s}') \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + 1) \\
&= \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + 1).
\end{aligned}$$

Summing up over all players $i \in [n]$ yields

$$\begin{aligned}
\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P}) &= \sum_{i \in [n]} \sum_{s \in S} p(\mathsf{s}) \sum_{e \in s_i} f_e(\delta_e(\mathsf{s})) \\
&\leq \sum_{i \in [n]} \sum_{s \in S} p(\mathsf{s}) \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + 1) \\
&= \sum_{s \in S} p(\mathsf{s}) \sum_{e \in E} \delta_e(\mathbf{Q}) \cdot f_e(\delta_e(\mathsf{s}) + 1).
\end{aligned}$$

Now, $\delta_e(\mathbf{Q})$ and $\delta_e(\mathsf{s})$ are both integer, since $\mathbf{Q}$ and s are both pure strategy profiles. Thus, by choosing c_1, c_2 such that

$$y \cdot f(x+1) \leq c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y) \tag{5.3}$$

for all polynomials f with maximum degree d and non-negative coefficients and for all $x, y \in \mathbb{N}_0$, we get

$$\begin{aligned}\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P}) &\le \sum_{s\in S} p(\mathsf{s}) \sum_{e\in E} \left[c_1\cdot\delta_e(\mathsf{s})\cdot f_e(\delta_e(\mathsf{s})) + c_2\cdot\delta_e(\mathbf{Q})\cdot f_e(\delta_e(\mathbf{Q}))\right] \\ &= c_1\cdot \mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P}) + c_2\cdot \mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{Q}).\end{aligned}$$

With $c_1 < 1$ it follows that $\frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P})}{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{Q})} \le \frac{c_2}{1-c_1}$. Since $\mathbf{P}$ is an arbitrary (mixed) Nash equilibrium we get

$$\mathsf{PoA}_{\mathsf{TL}} \le \frac{c_2}{1-c_1}. \tag{5.4}$$

In fact, c_1 and c_2 depend on the maximum degree d, however, for the sake of readability we omit this dependence in our notation.

We will now show how to determine constants c_1 and c_2 such that Inequality (5.3) holds and such that the resulting upper bound of $\frac{c_2}{1-c_1}$ is minimal. To do so, we will first show, that it suffices to consider Inequality (5.3) with $y = 1$ and $f(x) = x^d$.

Since f is a polynomial of maximum degree d with non-negative coefficients, it is sufficient to determine c_1 and c_2 that fulfill (5.3) for $f(x) = x^r$ for all integers $0 \le r \le d$.

So let $f(x) = x^r$ for some $0 \le r \le d$. In this case (5.3) reduces to

$$y\cdot(x+1)^r \le c_1\cdot x^{r+1} + c_2\cdot y^{r+1}. \tag{5.5}$$

For any given constant $0 \le c_1 < 1$ let $c_2(r, c_1)$ be the minimum value for c_2 such that (5.5) holds, that is

$$\begin{aligned}c_2(r,c_1) &= \max_{x\in\mathbb{N}_0, y\in\mathbb{N}} \left\{\frac{y(x+1)^r - c_1\cdot x^{r+1}}{y^{r+1}}\right\} \\ &= \max_{x\in\mathbb{N}_0, y\in\mathbb{N}} \left\{\left(\frac{x+1}{y}\right)^r - c_1\cdot\left(\frac{x}{y}\right)^{r+1}\right\}.\end{aligned}$$

Note that (5.5) holds for any c_2 when $y = 0$. By Lemma 5.1 we have

$$c_2(r,c_1) = \max_{x\in\mathbb{N}_0} \left\{(x+1)^r - c_1\cdot x^{r+1}\right\}. \tag{5.6}$$

Now, $c_2(r, c_1)$ is the maximum of infinitely many linear functions in c_1; one for each $x \in \mathbb{N}_0$. Denote $\mathcal{F}_r$ as the (infinite) set of linear functions defining $c_2(r, c_1)$. Thus,

$$\mathcal{F}_r = \{g_x^{(r)} : (0,1) \to \mathbb{R} \mid g_x^{(r)}(c_1) = (x+1)^r - c_1\cdot x^{r+1}, x\in\mathbb{N}_0\}.$$

For the partial derivative of any function $(x, r, c_1) \mapsto g_x^{(r)}(c_1)$ we get

$$\begin{aligned}\frac{\partial((x+1)^r - c_1\cdot x^{r+1})}{\partial r} &= (x+1)^r\cdot\ln(x+1) - c_1\cdot x^{r+1}\cdot\ln(x) \\ &> \ln(x+1)\left[(x+1)^r - c_1\cdot x^{r+1}\right] \\ &\ge 0,\end{aligned}$$

for $(x+1)^r - c_1 \cdot x^{r+1} \geq 0$, that is, for the positive range of the chosen function from $\mathcal{F}_r$. Thus, the positive range of $(x+1)^d - c_1 \cdot x^{d+1}$ dominates the positive range of $(x+1)^r - c_1 \cdot x^{r+1}$ for all $0 \leq r \leq d$. Since $c_2(r, c_1) > 0$ for all $0 \leq r \leq d$, it follows that $c_2(d, c_1) \geq c_2(r, c_1)$, for all $0 \leq r \leq d$. Thus, without loss of generality, we may assume that $f(x) = x^d$.

For $s, t \in \mathbb{R}_{\geq 0}$ and $s \neq t$ define $\gamma(s,t)$ as the intersection abscissa of $g_s^{(d)}$ and $g_t^{(d)}$ (as in Lemma 5.2). Now consider the intersection of the two functions $g_v^{(d)}$ and $g_{v+1}^{(d)}$ from $\mathcal{F}_d$ for some $v \in \mathbb{N}$. We show that this intersection lies above all other functions from $\mathcal{F}_d$.

- First consider any function $g_z^{(d)}$ with $z > v+1$. We have $g_z^{(d)}(0) > g_{v+1}^{(d)}(0) > g_v^{(d)}(0)$. Furthermore, by Lemma 5.2 we get $\gamma(v,z) < \gamma(v, v+1)$. It follows that $g_v^{(d)}(\gamma(v, v+1)) > g_z^{(d)}(\gamma(v, v+1))$.
- Now consider any function $g_z^{(d)}$ with $z < v$. We have $g_{v+1}^{(d)}(0) > g_v^{(d)}(0) > g_z^{(d)}(0)$. Furthermore, by Lemma 5.2 we get $\gamma(v,z) > \gamma(v, v+1)$. Again, it follows that $g_v^{(d)}(\gamma(v, v+1)) > g_z^{(d)}(\gamma(v, v+1))$.

Thus, all intersections of two consecutive linear functions from $\mathcal{F}_d$ lie on $c_2(d, c_1)$. The structure of function $c_2(d, c_1)$ is illustrated in Figure 5.1.

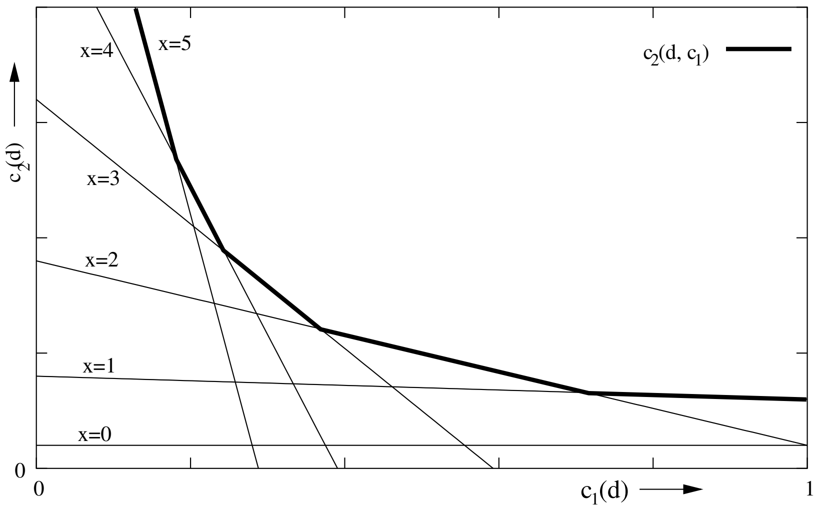

Fig. 5.1. The function $c_2(d, c_1)$

By (5.4), any point that lies on $c_2(d, c_1)$ gives an upper bound on $\mathsf{PoA}_{\mathsf{TL}}$. Let k be the largest integer such that $(k+1)^d \geq k^{d+1}$, that is $k = \lfloor \Phi_d \rfloor$. Then $(k+2)^d < (k+1)^{d+1}$. Choose c_1 and c_2 at the intersection of the two lines from $\mathcal{F}_d$ with $x = k$ and $x = k+1$, that is $c_2 = (k+1)^d - c_1 \cdot k^{d+1}$ and $c_2 = (k+2)^d - c_1 \cdot (k+1)^{d+1}$. Thus,

$$c_1 = \frac{(k+2)^d - (k+1)^d}{(k+1)^{d+1} - k^{d+1}} \quad \text{and} \quad c_2 = \frac{(k+1)^{2d+1} - (k+2)^d \cdot k^{d+1}}{(k+1)^{d+1} - k^{d+1}}.$$

Note that by the choice of k we have $0 < c_1 < 1$.

It follows that

$$\mathsf{PoA}_{\mathsf{TL}} \leq \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}.$$

This completes the proof of the theorem. ∎

5.2.2 Lower Bound

In Theorem 5.4 we give a matching lower bound which also holds for unweighted network congestion games (Corollary 5.5).

Theorem 5.4. *For unweighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have*

$$\mathsf{PoA}_{\mathsf{TL}} \geq \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$.

Proof. Given the maximum degree $d \in \mathbb{N}$ for the polynomial latency functions, we construct a congestion game for $n \geq k+2$ players and $|E| = 2n$ facilities.

We divide the set E into two subsets $E_1 = \{g_1, \ldots, g_n\}$ and $E_2 = \{h_1, \ldots, h_n\}$. Each player i has two pure strategies, $P_i = \{g_{i+1}, \ldots, g_{i+k}, h_{i+1}, \ldots, h_{i+k+1}\}$ and $Q_i = \{g_i, h_i\}$ where $g_j = g_{j-n}$ and $h_j = h_{j-n}$ for $j > n$. I. e. $S_i = \{Q_i, P_i\}$.

Each of the facilities in E_1 share the latency function $x \mapsto ax^d$ for an $a \in \mathbb{R}_{>0}$ (yet to be determined) whereas the facilities in E_2 have latency $x \mapsto x^d$.

Obviously, the optimal allocation $\mathbf{Q}$ is for every player i to choose Q_i. Now we determine a value for a such that the allocation $\mathbf{P} = (P_1, \ldots, P_n)$ becomes a Nash equilibrium, i. e. each player i is satisfied with $\mathbf{P}$, that is $\mathsf{PC}_i(\mathbf{P}) \leq \mathsf{PC}_i(\mathbf{P}_{-i}, Q_i)$ for all $i \in [n]$, or equivalently

$$k \cdot a \cdot k^d + (k+1) \cdot (k+1)^d \leq a \cdot (k+1)^d + (k+2)^d.$$

Resolving to the coefficient a gives

$$a \geq \frac{(k+1)^{d+1} - (k+2)^d}{(k+1)^d - k^{d+1}} > 0. \tag{5.7}$$

Because $(k+1)^d \neq k^{d+1}$ due to either $k+1$ or k being odd and the other being even, a is well defined and positive. Now since for any player i the private costs are $\mathsf{PC}_i(\mathbf{Q}) = a + 1$ and $\mathsf{PC}_i(\mathbf{P}) = a \cdot k^{d+1} + (k+1)^{d+1}$, it follows that

$$\frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P})}{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{Q})} = \frac{\sum_{i \in [n]} \mathsf{PC}_i(\mathbf{P})}{\sum_{i \in [n]} \mathsf{PC}_i(\mathbf{Q})} = \frac{a \cdot k^{d+1} + (k+1)^{d+1}}{a+1}. \tag{5.8}$$

Provided that $(k+1)^d \geq k^{d+1}$, it is not hard to see that (5.8) is monotonically decreasing in a. Thus, we assume equality in (5.7), which then gives

$$\mathsf{PoA}_{\mathsf{TL}} \geq \frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P})}{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{Q})} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}.$$

This completes the proof of the theorem. ■

We close this section by showing that the just shown lower bound also holds for unweighted *network* congestion games.

Corollary 5.5. *The lower bound in Theorem 5.4 on* $\mathsf{PoA}_{\mathsf{TL}}$ *also holds for unweighted network congestion games.*

Proof. Instances of the congestion game in Theorem 5.4 can be characterized by two parameters: the maximum degree d of the latency functions and the number of players $n \geq \lfloor \Phi_d \rfloor + 2$. The number of edges is then given by $2n$.

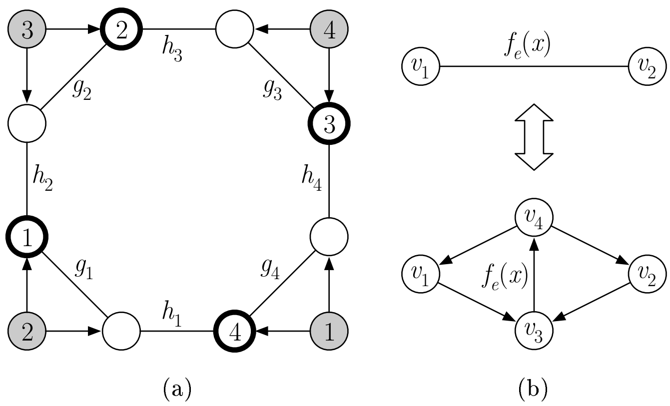

Fig. 5.2. Network congestion game for $d = 2$ and 4 players

Figure 5.2 (a) shows an example of the network congestion game for quadratic latency functions (i.e. $d = 2$) and for $n = 4$ players. Unlabeled edges have $f_e(x) = 0$ as their latency function. We say that these edges are *free*. All other edges have the associated latency function as in Theorem 5.4. In the following we outline the general construction scheme.

The network corresponding to an instance characterized by (d, n) can be constructed as follows:

There is a circle of $2n$ undirected edges $g_1, h_1, g_2, h_2, \ldots, g_n, h_n$. Each undirected edge (v_1, v_2) has to be replaced by the construction shown in Figure 5.2 (b). This insures that no matter in which direction a player uses edge (v_1, v_2) it produces load on the directed edge (v_3, v_4).

Now, every player i has its own origin node outside the circle – which is indicated by a gray background in the example. This node has an edge to the connecting node of g_i and h_{i-1}. The destination node of each player i is the node between h_i and g_{i+1}, represented by a thick outline in the figure. To also allow for the strategy P_i for each player i (which essentially goes the other way round inside the circle), we finally add one more edge from i's origin node to the connecting node of h_{i+k+1} and g_{i+k+1}. As before, let $k = \lfloor \Phi_d \rfloor$. ■

5.3 Price of Anarchy for Weighted Congestion Games

In this section, we prove the exact value for the price of anarchy of weighted congestion games with polynomial latency functions.

We prove the upper bound in Theorem 5.6. In Theorem 5.7 we give a matching lower bound which also holds for weighted network congestion games (Corollary 5.9). Corollary 5.10 shows the impact of player weights to the price of anarchy.

5.3.1 Upper Bound

Theorem 5.6. *For weighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients we have $\mathsf{PoA}_{\mathsf{TL}} \leq \Phi_d^{d+1}$.*

Proof. Let $\mathbf{P} = (P_1, \ldots, P_n)$ be a (mixed) Nash equilibrium and let $\mathbf{Q} = (Q_1, \ldots, Q_n)$ be a pure strategy profile with optimum total latency. We first note that due to the Nash inequalities it holds that

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{P}) &= \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in s_i} f_e(\delta_e(\mathsf{s})) \\
&\leq \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + w_i) \\
&\leq \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + \delta_e(\mathbf{Q})).
\end{aligned}$$

This gives the following upper bound for the total latency:

$$\begin{aligned}
\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P}) &= \sum_{i=1}^{n} w_i \cdot \mathsf{PC}_i(\mathbf{P}) \\
&\leq \sum_{i=1}^{n} w_i \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in Q_i} f_e(\delta_e(\mathsf{s}) + \delta_e(\mathbf{Q})) \\
&= \sum_{\mathsf{s} \in S} p(\mathsf{s}) \sum_{e \in E} \delta_e(\mathbf{Q}) \cdot f_e(\delta_e(\mathsf{s}) + \delta_e(\mathbf{Q})).
\end{aligned}$$

Similarly to Theorem 5.3, by choosing $c_1, c_2 \in \mathbb{R}$ such that

$$y \cdot f(x+y) \leq c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y) \tag{5.9}$$

for all polynomials f with maximum degree d and non-negative coefficients and for all $x, y \in \mathbb{R}_{\geq 0}$, we get

$$\begin{aligned}\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P}) &\leq \sum_{\mathsf{s}\in S} p(\mathsf{s}) \sum_{e\in E}\Big[c_1\cdot\delta_e(\mathsf{s})\cdot f_e(\delta_e(\mathsf{s})) + c_2\cdot\delta_e(\mathbf{Q})\cdot f_e(\delta_e(\mathbf{Q}))\Big]\\ &= c_1\cdot\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P}) + c_2\cdot\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{Q}).\end{aligned}$$

Note here that (5.9) varies from (5.3) of Theorem 5.3 and hence the former values for c_1 and c_2 in Theorem 5.3 cannot simply be reused.

For (5.9) to hold for all $x, y \in \mathbb{R}_{\geq 0}$, both c_1 and c_2 must be non-negative. With $0 \leq c_1 < 1$ it follows that $\frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{P})}{\mathsf{SC}_{\mathsf{TL}}(\Gamma,\mathbf{Q})} \leq \frac{c_2}{1-c_1}$. Since $\mathbf{P}$ is an arbitrary Nash equilibrium we get $\mathsf{PoA}_{\mathsf{TL}} \leq \frac{c_2}{1-c_1}$.

Now let $r \in [d]_0$. Because of the equivalences

$$\begin{aligned}\forall x,y\in\mathbb{R}_{\geq 0}: &\qquad y\cdot(x+y)^r \leq c_1\cdot x^{r+1} + c_2\cdot y^{r+1}\\ \Leftrightarrow \forall x\in\mathbb{R}_{\geq 0}, y\in\mathbb{R}_{>0}: &\qquad \left(\frac{x}{y}+1\right)^r \leq c_1\cdot\left(\frac{x}{y}\right)^{r+1} + c_2\\ \Leftrightarrow \forall x\in\mathbb{R}_{\geq 0}: &\qquad (x+1)^r \leq c_1\cdot x^{r+1} + c_2\end{aligned}$$

it is sufficient to choose c_2 depending on $c_1 \in (0,1)$ and $r \in [d]_0$ as

$$c_2 \geq c_2(r, c_1) = \max_{x\in\mathbb{R}_{\geq 0}} \{(x+1)^r - c_1\cdot x^{r+1}\},$$

in order to fulfill (5.9) for every monomial f of degree r. With the same argument as for (5.6) in Theorem 5.3 it is sufficient to consider only the monomial of the largest degree d, so that (5.9) will then hold for any polynomial of maximum degree d with positive coefficients.

Let again $\mathcal{F}_d$ denote the infinite set of linear functions defining $c_2(d, c_1)$, i. e.

$$\mathcal{F}_d = \{g_x^{(d)} : (0,1) \to \mathbb{R} \mid g_x^{(d)}(c_1) = (x+1)^d - c_1\cdot x^{d+1}, x\in\mathbb{R}_{\geq 0}\}.$$

Keep d fixed and – like in Lemma 5.2 – define $\gamma(s,t)$ for $s,t \in \mathbb{R}_{\geq 0}$ and $s \neq t$ as the intersection abscissa of $g_s^{(d)}$ and $g_t^{(d)}$. From Lemma 5.2, we know that $x \mapsto \gamma(v,x)$ is both continuous and strictly decreasing on $(0,v)$ and then again on (v,∞), for any $v \in \mathbb{R}_{>0}$. We are interested in the limit of $x \mapsto \gamma(v,x)$ at $x = v$ and get

$$\lim_{\varepsilon\to 0}\gamma(v, v+\varepsilon) = \lim_{\varepsilon\to 0}\frac{(v+1+\varepsilon)^d - (v+1)^d}{(v+\varepsilon)^{d+1} - v^{d+1}} = \lim_{\varepsilon\to 0}\frac{\sum_{i=0}^{d-1}\binom{d}{i}\cdot(v+1)^i\cdot\varepsilon^{d-i}}{\sum_{i=0}^{d}\binom{d+1}{i}\cdot v^i\cdot\varepsilon^{d+1-i}}$$

which yields (by canceling one ε)

$$\lim_{\varepsilon\to 0}\gamma(v, v+\varepsilon) = \frac{d\cdot(v+1)^{d-1}}{(d+1)\cdot v^d}. \qquad (5.10)$$

Note here that this limit exists regardless of the direction in which ε approaches 0. Therefore, we make the natural extension of defining $\gamma(v,v)$ to be the limit from (5.10). We observe that this extension makes $x \mapsto \gamma(v,x)$ strictly decreasing and continuous in all of its domain $\mathbb{R}_{>0}$.

Now consider the intersection of the two functions $g_v^{(d)}$ and $g_{v+\varepsilon}^{(d)}$ from $\mathcal{F}_d$ for $\varepsilon \to 0$ and for some $v \in \mathbb{R}_{>0}$. For $\varepsilon \to 0$ these functions intersect at $(\gamma(v,v), g_v^{(d)}(\gamma(v,v)))$. We show that this intersection lies above all other functions from $\mathcal{F}_d$.

- First consider any function $g_z^{(d)}$ with $z > v$. We have $g_z^{(d)}(0) > g_{v+\varepsilon}^{(d)}(0) > g_v^{(d)}(0)$ for $0 < \varepsilon < z - v$. Furthermore, by Lemma 5.2 we get $\gamma(v,z) < \gamma(v, v+\varepsilon) < \gamma(v,v)$. It follows that $g_v^{(d)}(\gamma(v,v)) > g_z^{(d)}(\gamma(v,v))$.
- Now consider any function $g_z^{(d)}$ with $z < v$. We have $g_{v+\varepsilon}^{(d)}(0) > g_v^{(d)}(0) > g_z^{(d)}(0)$. Furthermore, by Lemma 5.2 we get $\gamma(v,z) > \gamma(v,v) > \gamma(v,v+\varepsilon)$. Again, it follows that $g_v^{(d)}(\gamma(v,v)) > g_z^{(d)}(\gamma(v,v))$.

Thus, $g_v^{(d)}(\gamma(v,v)) = \max_{x \in \mathbb{R}_{\geq 0}} g_x^{(d)}(\gamma(v,v))$ for all $v \in \mathbb{R}_{>0}$.

Therefore, we can express $c_2(d, c_1)$ as

$$c_2(d, c_1) = \max_{x \in \mathbb{R}_{\geq 0}} \{g_x^{(d)}(c_1)\} = g_s^{(d)}(c_1) \tag{5.11}$$

where $s \in \mathbb{R}_{\geq 0}$ solely depends on c_1 and d and is defined by $\gamma(s,s) = c_1$. That means $g_s^{(d)}$ and $g_{s+\varepsilon}^{(d)}$, for $\varepsilon \to 0$, have c_1 as their intersection abscissa. We choose

$$\begin{aligned} c_1 &= \gamma(\Phi_d, \Phi_d) \\ &= \frac{d \cdot (\Phi_d + 1)^{d-1}}{(d+1) \cdot \Phi_d^d} \\ &= \frac{d \cdot \Phi_d^{d+1}}{(d+1) \cdot (\Phi_d + 1) \cdot \Phi_d^d} \\ &= \frac{d \cdot \Phi_d}{(d+1)(\Phi_d + 1)} \in (0,1). \end{aligned}$$

With (5.11), this yields $c_2(d, c_1) = g_{\Phi_d}^{(d)}(c_1) = \Phi_d^{d+1} \cdot (1 - c_1)$. Thus, we can choose $c_2 = \Phi_d^{d+1} \cdot (1 - c_1)$ and get

$$\mathsf{PoA}_{\mathsf{TL}} \leq \frac{c_2}{1 - c_1} = \frac{\Phi_d^{d+1} \cdot (1 - c_1)}{1 - c_1} = \Phi_d^{d+1}.$$

This completes the proof of the theorem. ∎

5.3.2 Lower Bound

Theorem 5.7. *For weighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have $\mathsf{PoA}_{\mathsf{TL}} \geq \Phi_d^{d+1}$.*

Proof. Given the maximum degree $d \in \mathbb{N}$ for the polynomial latency functions, set $k \geq \max\{\binom{d}{\lfloor d/2 \rfloor}, 2\}$. Note, that $\binom{d}{\lfloor d/2 \rfloor} = \max_{j \in [d]_0} \binom{d}{j}$. We construct a congestion game for $n = (d+1) \cdot k$ players and $|E| = n$ facilities.

We divide the set E into $d+1$ partitions:
For $i \in [d]_0$, let $E_i = \{g_{i,1}, \ldots, g_{i,k}\}$, with each $g_{i,j}$ sharing the latency function $x \mapsto a_i \cdot x^d$. The values of the coefficients a_i will be determined later. For simplicity of notation, set $g_{i,j} = g_{i,j-k}$ for $j > k$ in the following.

Similarly, we partition the set of players $[n]$:
For $i \in [d]_0$, let $N_i = \{u_{i,1}, \ldots, u_{i,k}\}$. The weight of each player in set N_i is Φ_d^i, so $w_{u_{i,j}} = \Phi_d^i$ for all $i \in [d]_0$, $j \in [k]$.

Now, for every set N_i, each player $u_{i,j} \in N_i$ has exactly two strategies:

$$Q_{u_{i,j}} = \{g_{i,j}\} \quad \text{and} \quad P_{u_{i,j}} = \begin{cases} \{g_{d,j+1}, \ldots, g_{d,j+\binom{d}{i}}, g_{i-1,j}\} & \text{for } i = 1 \text{ to } d, \\ \{g_{d,j+1}\} & \text{for } i = 0. \end{cases}$$

Now let $\mathbf{Q} = (Q_1, \ldots, Q_n)$ and $\mathbf{P} = (P_1, \ldots, P_n)$ be strategy profiles. The facilities in each set E_i then have the following loads for $\mathbf{Q}$ and $\mathbf{P}$:

	load on every facility $e \in E_i$	
i	$\delta_e(\mathbf{Q})$	$\delta_e(\mathbf{P})$
d	Φ_d^d	$\sum_{l=0}^{d} \binom{d}{l} \Phi_d^l = (\Phi_d + 1)^d = \Phi_d^{d+1}$
0 to $d-1$	Φ_d^i	Φ_d^{i+1}

For $\mathbf{P}$ to become a Nash Equilibrium, we need to fulfill the following Nash inequalities for each set N_i of players:

i	Nash inequality to fulfill
1 to d	$\mathsf{PC}_{u_{i,j}}(\mathbf{P}) = \binom{d}{i} \cdot a_d \cdot (\Phi_d^{d+1})^d + a_{i-1} \cdot (\Phi_d^i)^d \le a_i \cdot (\Phi_d^{i+1} + \Phi_d^i)^d = \mathsf{PC}_{u_{i,j}}(\mathbf{P}_{-u_{i,j}}, Q_{u_{i,j}})$
0	$\mathsf{PC}_{u_{0,j}}(\mathbf{P}) = a_d \cdot (\Phi_d^{d+1})^d \le a_0 \cdot (\Phi_d + 1)^d = \mathsf{PC}_{u_{0,j}}(\mathbf{P}_{-u_{0,j}}, Q_{u_{0,j}})$

Replacing "$\le$" by "$=$" yields a homogeneous system of linear equations, i. e. the system $B_d \cdot a = 0$ where B_d is the following $(d+1) \times (d+1)$ matrix:

$$B_d = \begin{pmatrix}
-\Phi_d^{d^2+d+1} + \Phi_d^{d^2+d} & \Phi_d^{d^2} & 0 & \cdots & & & & \cdots & 0 \\
\binom{d}{d-1}\Phi_d^{d^2+d} & -\Phi_d^{d^2+1} & \ddots & & & & & & \vdots \\
\vdots & 0 & \ddots & & & & & & \\
\vdots & \vdots & & \ddots & \ddots & & & & \vdots \\
\binom{d}{i}\Phi_d^{d^2+d} & 0 & \cdots & 0 & -\Phi_d^{id+d+1} & \Phi_d^{id} & 0 & \cdots & 0 \\
\vdots & \vdots & & & 0 & \ddots & \ddots & & \vdots \\
 & & & & \vdots & & & \ddots & 0 \\
\vdots & \vdots & & & \vdots & & & \ddots & \Phi_d^d \\
\Phi_d^{d^2+d} & 0 & \cdots & & 0 & \cdots & & 0 & -\Phi_d^{d+1}
\end{pmatrix} \tag{5.12}$$

and $a = (a_d \dots a_0)^t$. Obviously, a solution to this system fulfills the initial Nash inequalities. Note, that

$$(\Phi_d^{i+1} + \Phi_d^i)^d = (\Phi_d^i)^d \cdot (\Phi_d + 1)^d = \Phi_d^{id+d+1}.$$

Claim 5.8. The $(d+1) \times (d+1)$ matrix B_d from (5.12) has rank d.

Proof. We use the well known fact from linear algebra that if a matrix C results from another matrix D by adding a multiple of one row (or column) to another row (or column, respectively) then $\mathrm{rank}(C) = \mathrm{rank}(D)$.

Now consider the matrix C_d that results from adding row j multiplied by the factor Φ_d^{-1} to row $j-1$, sequentially done for $j = d+1, d, \dots, 2$. Obviously, C_d is a lower triangular matrix with nonzero elements only in the first column and on the principal diagonal.

For the top left element of C_d we get

$$-\Phi_d^{d^2+d+1} + \sum_{j=0}^{d} \binom{d}{j} \Phi_d^{d^2+j} = \Phi_d^{d^2} \cdot \Bigg(-\Phi_d^{d+1} + \underbrace{\sum_{j=0}^{d} \binom{d}{j} \Phi_d^j}_{(\Phi_d+1)^d} \Bigg) = 0.$$

Since all elements on the principal diagonal of C_d – with the just shown exception of the first one – are nonzero, it is easy to see that C_d (and thus also B_d) has rank d. ∎

By the above claim it follows that the column vectors of B_d are linearly dependent and thus there are – with degree of freedom 1 – infinitely many linear combinations of them yielding 0. In other words, $B_d \cdot a = 0$ has a one-dimensional solution space.

We now show (by induction over i) that all coefficients a_i, $i \in [d]_0$ must have the same sign and thus we can alway find a valid solution. From the last equality, for $i = 0$, we have that a_d and a_0 must have the same sign. Now for $i = 1, \dots, d-1$, it follows that a_i must have the same sign as a_{i-1} and a_d, for $(\Phi_d^{d+1})^d$, $(\Phi_d^i)^d$, and $(\Phi_d^{i+1} + \Phi_d^i)^d$ are all positive.

Choosing $a \neq 0$ with all components being positive, all coefficients of the latency functions are positive. We get,

$$\mathsf{PoA}_{\mathsf{TL}} \geq \frac{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{P})}{\mathsf{SC}_{\mathsf{TL}}(\Gamma, \mathbf{Q})} = \frac{k \cdot \sum_{i=0}^{d} a_i (\Phi_d^{i+1})^{d+1}}{k \cdot \sum_{i=0}^{d} a_i (\Phi_d^{i})^{d+1}} = \Phi_d^{d+1}.$$

∎

We proceed to show that just given lower bound also holds for weighted *network* congestion games.

Corollary 5.9. *The lower bound in Theorem 5.7 on* $\mathsf{PoA}_{\mathsf{TL}}$ *also holds for weighted network congestion games.*

Proof. Each instance of the congestion game in Theorem 5.7 essentially can be characterized by two parameters: The maximum degree d of the latency functions and the number of facilities $k \geq \max\{\binom{d}{\lfloor d/2 \rfloor}, 2\}$ in each class E_i, where $i \in [d]_0$. Remember that the number of facilities – as well as the number of players – is given by $(d+1) \cdot k$.

Confer Figure 5.3 for an example in the case of quadratic latency functions (i. e. $d = 2$) and $k = 2$. For their respective players, gray nodes denote origins, whereas nodes with a thick outline represent destinations. Note that for the sake of clarity not all edges are shown, as will be explained later. Edges without a label have $f_e(x) = 0$ as their latency function. Again, we call these edges *free* edges. All other edges have the associated latency function as in Theorem 5.7. In the following, we will outline the general construction scheme.

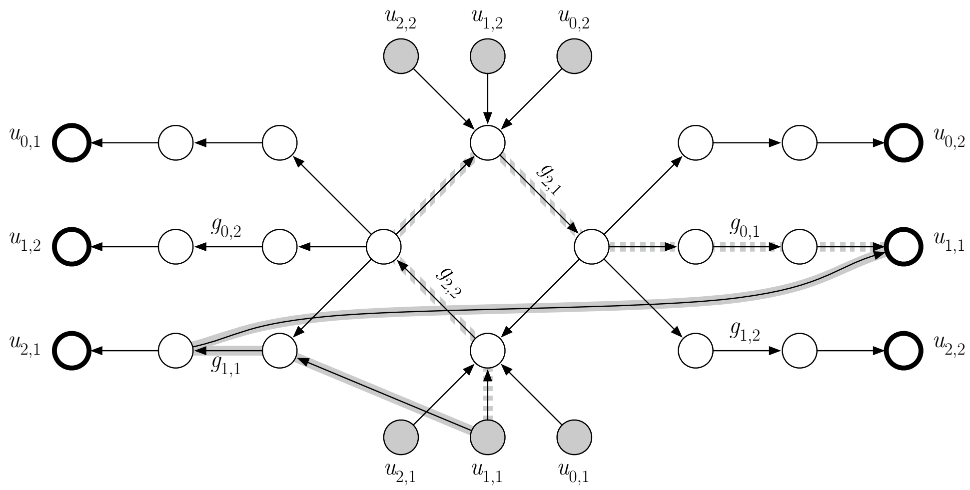

Fig. 5.3. Network congestion game for $d = 2$ and $k = 2$

The network corresponding to an instance characterized by (d, k) can be constructed as follows:

There is a circle of $2 \cdot k$ edges where every other edge represents a resource $g_{d,1}, g_{d,2}, \ldots, g_{d,k}$. All remaining edges in the circle are free edges. Furthermore, every player $u_{i,j}$ has its own origin node which has a single free edge to $g_{d,j+1}$. Consequently, circle edge $g_{d,j+\binom{d}{i}}$ connects to a free edge which then in turn connects to edge $g_{i-1,j}$. (In case $i = 0$, the latter simply is another free edge.) From there, there is another free edge to the destination node of player $u_{i,j}$. Note that, thus far, the graph has exactly one acyclic path for each player, i. e. for each origin-destination pair. Each of these paths represents that player's "unfavorable" strategy which has been denoted as $P_{u_{i,j}}$ in Theorem 5.7.

One can now add two more free edges, for each player $u_{i,j}$, that allow $u_{i,j}$ to also use its optimal strategy $Q_{u_{i,j}}$: From $u_{i,j}$'s origin node add a free link to $g_{i,j}$, and from $g_{i,j}$ add a free link to $u_{i,j}$'s destination node. We call the first type of links *A-Links*, the latter *B-Links*. Note that in Figure 5.3, A- and B-Links are only shown for player $u_{1,1}$. (The figure is complete otherwise.) The thick gray path denotes player $u_{1,1}$'s strategy in the system optimum, whereas the hatched path indicates its strategy in the worst-case Nash equilibrium.

A-Links obviously cannot create shortcuts for *other* players as origin nodes only have outgoing edges. Similarly, destination nodes only have incoming edges and therefore B-Links cannot create shortcuts for other players, either. Eventually, neither A- nor B-Links can create a shortcut for the same player's other strategy $P_{u_{*,*}}$ as they do not share any nodes, except for the origin and destination nodes.

Note, however, that B-Links do create additional paths: In Figure 5.3, for instance, player $u_{1,1}$ now has the further option of using a path consisting of five edges: three free ones, $g_{2,2}$, and $g_{1,1}$. Nevertheless, all such additional paths are supersets of the player's optimal strategy and thus neither change the system optimum nor the worst-case Nash equilibrium. ■

We close this section by studying the impact of weights to the price of anarchy.

Corollary 5.10. *The exact price of anarchy for* unweighted congestion games

$$\mathsf{PoA}_{\mathsf{TL}} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$*, is bounded by* $\lfloor \Phi_d \rfloor^{d+1} \leq \mathsf{PoA}_{\mathsf{TL}} \leq \Phi_d^{d+1}$.

Proof. The upper bound obviously is a direct consequence of Theorem 5.6. For the lower bound, define $A = (k+1)^d - k^{d+1}$ and $B = (k+1)^{d+1} - (k+2)^d$. Note that $A, B > 0$ by choice of k. Then,

$$\begin{aligned}\frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}} &= \frac{(A + k^{d+1})(k+1)^{d+1} - k^{d+1}(k+2)^d}{A+B} \\ &= k^{d+1} \cdot \frac{A \cdot \left(\frac{k+1}{k}\right)^{d+1} + B}{A+B} \geq \lfloor \Phi_d \rfloor^{d+1}.\end{aligned}$$

This completes the proof of the lower bound. ■

5.4 Conclusion and Discussion

In this chapter we closed the problem of obtaining the exact value of the price of anarchy for (weighted) congestion games with polynomial latency functions. We assumed polynomial latency functions of maximum degree d with non-negative coefficients. We considered the cases of unweighted and weighted players. The given results improve on the two recent STOC papers of Awerbuch et al. [5] and Christodoulou and Koutsoupias [17]. Our bounds on the price of anarchy depend on a new combinatorial sequence Φ_d, which is a generalization of the golden ratio to higher dimensions. We believe that this sequence is of independent combinatorial interest.

The key to our results was an improved analysis of an optimization process that was also considered by Christodoulou and Koutsoupias [17]. Our improved analysis uses some new and non-trivial ideas.

The techniques used for proving our upper bounds are also practical for other models. In particular, by applying them, we were able to prove (tight) upper bounds on the price of anarchy for weighted congestion games and Wardrop games, both with player-specific linear latency functions [50]. Moreover, we deployed them to show upper bounds on the price of anarchy for Wardrop games with certain polynomial latency functions [24]. In both works [24, 50] we also used the total latency as our social cost measure.

6

Bayesian Routing Games

6.1 Introduction

In recent years, motivated by non-cooperative systems like the Internet, combining ideas from Game Theory and Theoretical Computer Science has become more and more attractive. In many of these large-scale, non-cooperative systems, users have only incomplete information about the system for several reasons. In his honored work, Harsanyi [58] introduced an elegant approach to studying non-cooperative games with *incomplete information*, where the players are uncertain about some parameters. To model such games, Harsanyi introduced the *Harsanyi transformation*, which converts a (strategic) game with incomplete information to a strategic 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* [12]. In Bayesian Decision Theory probabilities are used as a measure of the degree of belief a person has in some proposition.

In this chapter, we introduce a particular selfish routing game with incomplete information that we call *Bayesian routing game*. These games have been formally introduced in Section 3.3. Bayesian routing games are a generalization to the routing games on parallel links studied in Chapter 4. In a Bayesian routing game, each of n selfish *players* wishes to assign its *traffic* to one of m parallel *links*. Each link has a certain *capacity*, which specifies the rate at which the link processes traffic. In the case of *identical links*, all links have equal capacity. Link capacities vary arbitrary, in the case of *related links*. The *latency* of a link is the total traffic on the link divided by the capacity of the link. Players do not know each other's traffic. Following Harsanyi's approach, we introduce for each player a set of possible *types*. We assume that all *type sets* are finite. Each type of a player corresponds to some traffic. Furthermore, we assume that there is a joint probability distribution $\mathbf{\Psi}$, called *type distribution*, over the set of all possible type realizations. In general, $\mathbf{\Psi}$ can be arbitrary; however, sometimes we assume $\mathbf{\Psi}$ to be *independent* – in that case, $\mathbf{\Psi}$ is expressed as the product of n independent probability distributions, one for each player type set.

In a *pure strategy*, a player chooses for each of its types a particular link. So, a pure strategy is a function from the type set of a player to the set of links. In a *mixed strategy*, a player uses a probability distribution over all its possible pure strategies. A *strategy profile* specifies a strategy for each of the players. Users choose strategies in order to minimize their *private cost*, which is defined as the expected latency experienced by the player. Note that due to the Bayesian model, the private cost in a pure strategy profile is given by the expectation over the type distribution $\mathbf{\Psi}$. For mixed strategy profiles, the expectation is taken over both the type distribution $\mathbf{\Psi}$ and the mixed strategies of the players.

The players neither cooperate with each other nor adhere to a global objective function, the so called *social cost* [67]. A stable state in which no player has an incentive to unilaterally change its strategy is called a *Bayesian Nash equilibrium*. In our study, we distinguish between *pure* and *mixed* Bayesian Nash equilibria. Of special interest to our work are *fully mixed* Bayesian Nash equilibria, where each player assigns strictly positive probability to each of its pure strategies.

If each player has only a single type, so that players are completely informed about each other's traffic, then we are in the setting of the routing games on parallel links (with complete information) that we studied in Chapter 4. In the following, we call them *complete information routing games*, in order to emphasize the connection to Bayesian routing games. In this setting, Bayesian Nash equilibria become Nash equilibria.

As before, we use the *price of anarchy* [86], as a measure of the maximum performance degradation due to the selfish behavior of the players. The price of anarchy can be defined with respect to different social cost measures.

As a generalization of our Bayesian routing games we also introduce *weighted Bayesian congestion games*. Weighted Bayesian congestion games generalize the weighted congestion games (as studied in Chapter 5) by incorporating incomplete information about the players' traffic. So, in a weighted Bayesian congestion game, the strategy set of each player is a subset of the power set of given resources. Weighted Bayesian congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem where the players do not know each other's traffic.

6.1.1 Summary of Results

Due to the new dimension that the incomplete information introduces to the routing game, 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:

6.1.1.1 Existence and Computational Complexity of Pure Bayesian Nash Equilibria

We define a new potential function that we use to prove that every weighted Bayesian congestion game possesses a pure Bayesian Nash equilibrium (Theorem 6.1). Observe that this existence result applies for the class of *weighted Bayesian congestion games*.

For the case of Bayesian routing games, identical links and independent type distributions, a pure Bayesian Nash equilibrium can be computed in polynomial time (Theorem 6.2). This computation is based on Graham's LPT scheduling algorithm [56]. 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.

6.1.1.2 Properties of Fully Mixed Bayesian Nash Equilibria

We show that for the case of identical links, the private cost of each player is maximized in a fully mixed Bayesian Nash equilibrium (Theorem 6.7). 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 might exist more than one fully mixed Bayesian Nash equilibrium, and we study their structural properties (Theorem 6.9). Finally, we determine the dimension of the space of fully mixed Bayesian Nash equilibria for the case of independent type distributions (Theorem 6.10).

6.1.1.3 Bounds on the Price of Anarchy

We close this chapter with bounds on the price of anarchy for three different social cost measures and for the case of identical links.

- The *makespan social cost*, which is defined by the expected maximum latency on a link, is a social cost measure that expresses the social welfare of the system. Here, we are able to show lower and upper bounds on the price of anarchy for different special cases (Theorem 6.12, Theorem 6.15 and Theorem 6.16). The exact price of anarchy for this social cost measure remains open, even for the case of identical links.
- A social cost measure that describes average player welfare is the *sum of private costs*. In this setting, it follows that for the case of identical links, each fully mixed Bayesian Nash equilibrium has maximum social cost (Theorem 6.17). Using this fact, we prove an upper bound of $\frac{m+n-1}{m}$ on the price of anarchy for the case of identical links (Theorem 6.18). We prove that this bound is asymptotically tight, already for complete information routing games.

- We also study social cost as *maximum of private costs*. For identical links, 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 (Theorem 6.20).

6.1.2 Related Work

Bayesian routing games and weighted Bayesian congestion games generalize the games studied in Chapter 4 and Chapter 5. So, many results that are already described in Section 4.1.2 and Section 5.1.2 are also of interest here. To keep this chapter self-contained, we again include those results that are most closely related.

Rosenthal [88] introduced the class of *congestion games* and showed that they always possess a pure Nash equilibrium. Fotakis et al. [38] considered *weighted congestion games* and proved the existence of a pure Nash equilibria, for the case where resources have linear latency functions. They also showed that a pure Nash equilibrium might not exist for weighted congestion games with general latency functions.

Harsanyi developed in his pioneering work [58, 59] a framework for studying competitive situations where the players have incomplete information. For an introduction to these so-called Bayesian games, we refer to Mas-Colell et al. [73] and Myerson [79]. Facchini et al. [28] considered Bayesian congestion models with players of identical weight, which have incomplete information about each other's preferences. Beier et al. [9] focused on a service provider congestion game with incomplete information.

Complete information routing games on parallel links were introduced by Koutsoupias and Papadimitriou [67]. Graham's LPT scheduling algorithm [56] computes a pure Nash equilibrium in this setting [37].

Mavronicolas and Spirakis [74] introduced the notion of *fully mixed Nash equilibria* to complete information routing games. They showed that, in case of existence, the fully mixed Nash equilibrium is unique. For the case of identical links, Gairing et al. [47] showed that fully mixed Nash equilibria maximize private costs.

For complete information routing games and makespan social cost, there exist tight bounds on the price of anarchy. These bounds are $\Theta(\frac{\log m}{\log\log m})$ for identical links [23, 66] and $\Theta(\frac{\log m}{\log\log\log m})$ for related links [23]. For complete information routing games and social cost defined as the sum of private cost, Berenbrink et al. [11] gave a lower bound of $\frac{n}{5}$ on the price of anarchy. Restricting to pure Nash equilibria they also showed an upper bound that solely depends on the players' traffic.

Subsequently to our work Georgiou et al. [53] introduced a routing game with incomplete information where the players have complete information about each other's traffic but only incomplete information about the latency functions in the network.

6.1.3 Organization

The rest of this chapter is organized as follows. Pure Bayesian Nash equilibria are studied in Section 6.2. Some interesting structural properties of fully mixed Bayesian Nash equilibria are treated in Section 6.3. Section 6.4 studies the price of anarchy. We conclude, in Section 6.5, with a summary of our results and some open problems.

6.2 Pure Bayesian Nash Equilibria

In this section we study the existence and the computational complexity of pure Bayesian Nash equilibria.

We first show that there is always a pure Bayesian Nash equilibrium in any Bayesian routing game. In fact, we show this result for the more general class of weighted Bayesian congestion games (Theorem 6.1). Then, we present a polynomial time algorithm, called PureBayesian, that computes a pure Bayesian Nash equilibrium for a Bayesian routing game with identical links and independent type distribution (Theorem 6.2). Finally, we show that PureBayesian cannot be used to compute a pure Bayesian Nash equilibrium for Bayesian routing games with *related* links (Proposition 6.3) or with *correlated* type distribution (Proposition 6.4).

6.2.1 Existence

We start by proving existence of pure Bayesian Nash equilibria.

Theorem 6.1. *Every weighted Bayesian congestion game Γ with linear latency functions has a pure Bayesian Nash equilibrium.*

Proof. Given a pure strategy profile $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$, define the function

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

We will prove that any unilateral strategy change of a type agent that decreases its private cost also decreases the value of the function Φ.
Given a pure strategy profile $\boldsymbol{\sigma}$, define for every player $r \in [n]$ and type $t \in T_r$,

$$\Phi_{(r,t)}(\boldsymbol{\sigma}) = \sum_{e\in\sigma_r(t)} \Psi(r,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi}|t_r = t)) + w(t)) + g_e(w(t))\right];$$

and for every resource $e \in [m]$ and player $r \in [n]$,

$$\Phi_e^{-r}(\boldsymbol{\sigma}) = \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\boldsymbol{\sigma}, (\boldsymbol{\Psi}|t_i = t)) + w(t)) + g_e(w(t))\right].$$

Observe first that

$$\begin{aligned}
&\sum_{t\in T_r} \Phi_{(r,t)}(\boldsymbol{\sigma}) + \sum_{e\in[m]} \Phi_e^{-r}(\boldsymbol{\sigma})\\
&= \sum_{t\in T_r}\sum_{e\in\sigma_r(t)} \Psi(r,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-r}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_r=t))+w(t))+g_e(w(t))\right]\\
&\quad + \sum_{e\in[m]}\sum_{i\in[n]\setminus\{r\}}\sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t))+w(t))+g_e(w(t))\right]\\
&= \sum_{t\in T_r}\sum_{e\in\sigma_r(t)} \Psi(r,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-r}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_r=t))+w(t))+g_e(w(t))\right]\\
&\quad + \sum_{i\in[n]\setminus\{r\}}\sum_{t\in T_i}\sum_{e\in\sigma_i(t)} \Psi(i,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t))+w(t))+g_e(w(t))\right]\\
&= \sum_{i\in[n]}\sum_{t\in T_i}\sum_{e\in\sigma_i(t)} \Psi(i,t)\cdot w(t)\cdot \left[g_e(\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t))+w(t))+g_e(w(t))\right]\\
&= \Phi(\boldsymbol{\sigma}).
\end{aligned}$$

Consider a unilateral strategy change of type agent $(r,\hat{t})$ from the set of resources $\sigma_r(\hat{t}) \in S_r$ to the set of resources $\sigma'_r(\hat{t}) \in S_r$. Set $\sigma'_r(t) = \sigma_r(t)$ for all $t \in T_r \setminus \{\hat{t}\}$ and define $\boldsymbol{\sigma'} = (\sigma_1,\ldots,\sigma_{r-1},\sigma'_r,\ \sigma_{r+1},\ldots\sigma_n)$ as the pure strategy profile resulting from $\boldsymbol{\sigma}$ after this strategy change. Assume that $v_{(r,\hat{t})}(\boldsymbol{\sigma'},\boldsymbol{\Psi}) < v_{(r,\hat{t})}(\boldsymbol{\sigma},\boldsymbol{\Psi})$, that is, the private cost of type agent $(r,\hat{t})$ decreases. Thus,

$$\begin{aligned}
&v_{(r,\hat{t})}(\boldsymbol{\sigma'},\boldsymbol{\Psi}) - v_{(r,\hat{t})}(\boldsymbol{\sigma},\boldsymbol{\Psi})\\
&= \sum_{e\in\sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\boldsymbol{\sigma'},(\boldsymbol{\Psi}|t_r=\hat{t}))+w(\hat{t})\right) - \sum_{e\in\sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_r=\hat{t}))+w(\hat{t})\right)\\
&< 0.
\end{aligned}$$

Moreover,

- $\Phi_{(r,t)}(\boldsymbol{\sigma}) = \Phi_{(r,t)}(\boldsymbol{\sigma'})$ for all type agents (r,t) where $t \in T_r \setminus \{\hat{t}\}$, and
- $\Phi_e^{-r}(\boldsymbol{\sigma}) = \Phi_e^{-r}(\boldsymbol{\sigma'})$ for all resources e that are neither in $\sigma_r(\hat{t})$ nor in $\sigma'_r(\hat{t})$ or that are in both $\sigma_r(\hat{t})$ as well as in $\sigma'_r(\hat{t})$, that is, $e \in \left([m] \setminus (\sigma_r(\hat{t}) \cup \sigma'_r(\hat{t}))\right) \cup \left(\sigma_r(\hat{t}) \cap \sigma'_r(\hat{t})\right)$. Observe that these are the resources where the load does not change.

Now, consider the change $\Delta(\Phi)$ to the function Φ due to this strategy change of type agent $(r,\hat{t})$. Clearly,

$$
\begin{aligned}
&\Delta(\Phi) \\
&= \Phi(\boldsymbol{\sigma'}) - \Phi(\boldsymbol{\sigma}) \\
&= \sum_{t \in T_r} \left(\Phi_{(r,t)}(\boldsymbol{\sigma'}) - \Phi_{(r,t)}(\boldsymbol{\sigma}))\right) + \sum_{e \in [m]} \left(\Phi_e^{-r}(\boldsymbol{\sigma'}) - \Phi_e^{-r}(\boldsymbol{\sigma})\right) \\
&= \left(\Phi_{(r,\hat{t})}(\boldsymbol{\sigma'}) - \Phi_{(r,\hat{t})}(\boldsymbol{\sigma}))\right) \\
&\quad + \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left(\Phi_e^{-r}(\boldsymbol{\sigma'}) - \Phi_e^{-r}(\boldsymbol{\sigma})\right) + \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left(\Phi_e^{-r}(\boldsymbol{\sigma'}) - \Phi_e^{-r}(\boldsymbol{\sigma})\right) \\
&= \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi),
\end{aligned}
$$

where

$$
\begin{aligned}
\Delta_1(\Phi) &= \Phi_{(r,\hat{t})}(\boldsymbol{\sigma'}) - \Phi_{(r,\hat{t})}(\boldsymbol{\sigma}), \\
\Delta_2(\Phi) &= \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left(\Phi_e^{-r}(\boldsymbol{\sigma'}) - \Phi_e^{-r}(\boldsymbol{\sigma})\right), \text{and} \\
\Delta_3(\Phi) &= \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left(\Phi_e^{-r}(\boldsymbol{\sigma'}) - \Phi_e^{-r}(\boldsymbol{\sigma})\right).
\end{aligned}
$$

Clearly,

$$
\begin{aligned}
&\Delta_1(\Phi) \\
&= \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(\sum_{e \in \sigma'_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma'}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right. \\
&\qquad \left. - \sum_{e \in \sigma_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right) \\
&= \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma'}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right. \\
&\qquad \left. - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right) \\
&= \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right. \\
&\qquad \left. - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\mathbf{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t}))\right] \right).
\end{aligned}
$$

Furthermore, due to the arrival of type agent $(r,\hat{t})$ on the resources $e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})$,

$$
\begin{aligned}
&\Delta_2(\Phi)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\\
&\qquad \cdot\Big(g_e(\delta_e^{-i}(\boldsymbol{\sigma}',(\boldsymbol{\Psi}|t_i=t))+w(t))+g_e(w(t))\\
&\qquad - g_e(\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t))+w(t))-g_e(w(t))\Big)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\\
&\qquad \cdot\Big(g_e(\delta_e^{-i}(\boldsymbol{\sigma}',(\boldsymbol{\Psi}|t_i=t)))-g_e(\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t)))\Big)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\\
&\qquad \cdot a_e\Big(\delta_e^{-i}(\boldsymbol{\sigma}',(\boldsymbol{\Psi}|t_i=t))-\delta_e^{-i}(\boldsymbol{\sigma},(\boldsymbol{\Psi}|t_i=t))\Big)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\cdot a_e\\
&\quad \cdot \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t}} \Psi(t_1,\dots,t_{i-1},t_{i+1},\dots,t_n|t_i=t)\cdot\left[\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma'_s(t_s)}} w(t_s)-\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma_s(t_s)}} w(t_s)\right].
\end{aligned}
$$

Note that

$$
\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma'_s(t_s)}} w(t_s)-\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma_s(t_s)}} w(t_s)=\begin{cases} w(\hat{t}), & \text{for all } (t_1,\dots,t_n)\in T \text{ where } t_r=\hat{t},\\ 0, & \text{else.}\end{cases}
$$

Hence,

$$
\begin{aligned}
&\Delta_2(\Phi)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \Psi(i,t)\cdot w(t)\\
&\qquad \cdot \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t,t_r=\hat{t}}} \Psi(t_1,\dots,t_{i-1},t_{i+1},\dots,t_n|t_i=t)\cdot w(\hat{t})\\
&= w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} w(t) \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t,t_r=\hat{t}}} \Psi(t_1,\dots,t_n)\\
&= w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t,t_r=\hat{t}}} \Psi(t_1,\dots,t_n)\cdot w(t)
\end{aligned}
$$

$$= \Psi(r,\hat{t}) \cdot w(\hat{t}) \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} a_e$$
$$\cdot \sum_{i \in [n] \setminus \{r\}} \sum_{\substack{t \in T_i: \\ e \in \sigma_i(t)}} \sum_{\substack{(t_1,\dots,t_n) \in T: \\ t_i = t, t_r = \hat{t}}} \Psi(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n | t_r = \hat{t}) \cdot w(t)$$
$$= \Psi(r,\hat{t}) \cdot w(\hat{t}) \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} a_e$$
$$\cdot \sum_{i \in [n] \setminus \{r\}} \sum_{\substack{(t_1,\dots,t_n) \in T: \\ e \in \sigma_i(t_i), t_r = \hat{t}}} \Psi(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n | t_r = \hat{t}) \cdot w(t_i)$$
$$= \Psi(r,\hat{t}) \cdot w(\hat{t}) \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} a_e$$
$$\cdot \sum_{\substack{(t_1,\dots,t_n) \in T: \\ t_r = \hat{t}}} \Psi(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n | t_r = \hat{t}) \sum_{\substack{i \in [n] \setminus \{r\}: \\ e \in \sigma_i(t_i)}} w(t_i)$$
$$= \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} a_e \cdot \delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})).$$

Similarly, since type agent $(r,\hat{t})$ left the resources $e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})$, we obtain that

$$\Delta_3(\Phi) = -\Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} a_e \cdot \delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})).$$

Hence,

$$\Delta(\Phi)$$
$$= \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi)$$
$$= \Psi(r,\hat{t}) \cdot w(\hat{t})$$
$$\cdot \left(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t})) + a_e \cdot \delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t}))\right] \right.$$
$$\left. - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t})) + a_e \cdot \delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t}))\right] \right)$$
$$= 2 \cdot \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) \right.$$
$$\left. - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) \right)$$
$$= 2 \cdot \Psi(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(\sum_{e \in \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi} | t_r = \hat{t})) + w(\hat{t})\right) \right.$$

$$- \sum_{e \in \sigma_r(\hat{t})} g_e \left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi}|t_r = \hat{t})) + w(\hat{t})\right)\Bigg)$$

$$= 2 \cdot \Psi(r, \hat{t}) \cdot w(\hat{t}) \cdot \Bigg(\sum_{e \in \sigma'_r(\hat{t})} g_e \left(\delta_e^{-r}(\boldsymbol{\sigma}', (\boldsymbol{\Psi}|t_r = \hat{t})) + w(\hat{t})\right)$$

$$- \sum_{e \in \sigma_r(\hat{t})} g_e \left(\delta_e^{-r}(\boldsymbol{\sigma}, (\boldsymbol{\Psi}|t_r = \hat{t})) + w(\hat{t})\right)\Bigg)$$

$$= 2 \cdot \Psi(r, \hat{t}) \cdot w(\hat{t}) \cdot \left(v_{(r,\hat{t})}(\boldsymbol{\sigma}', \boldsymbol{\Psi}) - v_{(r,\hat{t})}(\boldsymbol{\sigma}, \boldsymbol{\Psi})\right)$$
$$< 0 \,.$$

Thus, any unilateral strategy change of a type agent that decreases its private cost also decreases the value of the function Φ. Since the number of possible strategy profiles in Γ is finite, it follows that there is a pure strategy profile that minimizes Φ. In this strategy profile, no type agent can decrease its private cost by unilaterally changing its strategy. Hence, Γ has a pure Bayesian Nash equilibrium, as needed. ■

This generalizes a result by Fotakis et al. [38, Theorem 1] to the Bayesian setting. In particular our function Φ reduces to their potential function if each player has only a single type.

6.2.2 Computation

We now turn to the model of identical links and show how a pure Bayesian Nash equilibrium can be computed in polynomial time if the type distribution is independent. An algorithm, called PureBayesian, that performs this task is depicted in Figure 6.1. The algorithm computes a *normal* pure Bayesian Nash

---

PureBayesian(Γ)

Input: Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \boldsymbol{\Psi})$ with identical links and independent type distribution

Output: pure Bayesian Nash Equilibrium $\boldsymbol{\sigma}$

1: Calculate for each player $i \in [n]$ its expected traffic $W(i)$;

2: Construct a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t'_1, ..., t'_n)\}, 1)$ where $w(t'_i) = W(i)$ for all $i \in [n]$.

3: Compute a pure Nash equilibrium $\mathbf{L} = (\ell_1, \ldots, \ell_n)$ for Γ_{CI} in polynomial time with the LPT scheduling algorithm which assigns the players in order of non-increasing player traffic to minimum load links (see [37, 56]).

4: Set $\sigma_i(t) = \ell_i$ for all players $i \in [n]$ and types $t \in T_i$.

5: **return** $\boldsymbol{\sigma}$;

Fig. 6.1. PureBayesian

equilibrium, that is, for a fixed player, it assigns all types to the same link.

PUREBAYESIAN first computes for each player its expected traffic. Then, it uses Graham's LPT scheduling algorithm [56] to assign (all types of) the players in order non-increasing expected traffic to minimum load links. Gairing et al. [49] showed that the resulting strategy profile is a (normal) pure Bayesian Nash equilibrium.

Theorem 6.2 (Gairing et al. [49]). *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ 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 $\mathbf{\Psi}$ is represented in a compact form by a set of probabilities $\Psi(i,t)$ for $i \in [n]$ and $t \in T_i$.*

Algorithm PUREBAYESIAN 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 it always computes a *normal* pure Bayesian Nash equilibrium, while the following counter-examples show that a normal pure Bayesian Nash equilibrium does not exist in general.

Proposition 6.3. *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 $\Gamma = (2, 2, \mathbf{c}, T_1 \times T_2, \mathbf{\Psi})$ 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 $\Psi(1, t_1) = \Psi(1, t_1') = \frac{1}{2}$. We will now study the structure of pure Bayesian Nash equilibria for Γ and finally recognize that it has no *normal* pure Bayesian Nash equilibrium.

Let $\boldsymbol{\sigma}$ be an arbitrary pure Bayesian Nash equilibrium. Then,

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

while

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

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

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

So $\boldsymbol{\sigma}$ assigns t_1 to link 1 and t_1' to link 2. It follows that $\boldsymbol{\sigma}$ is the unique pure Bayesian Nash equilibrium. However, $\boldsymbol{\sigma}$ is not a *normal* pure Bayesian Nash equilibrium. The claim follows. ■

Proposition 6.4. *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 $\Gamma = (3, 2, \mathbf{1}, T_1 \times T_2 \times T_3, \mathbf{\Psi})$ 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 traffic $w(t_1) = w(t_2') = w(t_3) = w(t_3') = 1$ and $w(t_1') = w(t_2) = 2$. The correlated distribution $\mathbf{\Psi}$ is given by $\Psi(t_1, t_2, t_3) = \Psi(t_1', t_2', t_3') = \frac{1}{2}$.

Assume, by way of contradiction, that a *normal* pure Bayesian Nash equilibrium $\boldsymbol{\sigma}$ 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 and k be the two links. Without loss of generality, set $\sigma_1(t_1) = \sigma_1(t_1') = j$. Then, clearly

$$\lambda^j_{(2,t_2')}(\boldsymbol{\sigma}, \mathbf{\Psi}) \geq w(t_1') + w(t_2') = 3 \;\text{ while }\; \lambda^k_{(2,t_2')}(\boldsymbol{\sigma}, \mathbf{\Psi}) \leq w(t_3') + w(t_2') = 2.$$

Thus, $\sigma_2(t_2') = k$; hence, $\sigma_2(t_2) = \sigma_2(t_2') = k$ for all normal pure Bayesian Nash equilibria $\boldsymbol{\sigma}$. For the types of player 3, note that

$$\lambda^j_{(3,t_3)}(\boldsymbol{\sigma}, \mathbf{\Psi}) = w(t_1) + w(t_3) = 2 \;\text{ while }\; \lambda^k_{(3,t_3)}(\boldsymbol{\sigma}, \mathbf{\Psi}) = w(t_2) + w(t_3) = 3, \text{ and}$$
$$\lambda^j_{(3,t_3')}(\boldsymbol{\sigma}, \mathbf{\Psi}) = w(t_1') + w(t_3') = 3 \;\text{ while }\; \lambda^k_{(3,t_3')}(\boldsymbol{\sigma}, \mathbf{\Psi}) = w(t_2') + w(t_3') = 2.$$

Since $\boldsymbol{\sigma}$ is a Bayesian Nash equilibrium, $\sigma_3(t_3) = j$ and $\sigma_3(t_3') = k$. Hence, $\boldsymbol{\sigma}$ is not normal. A contradiction. ■

6.3 Properties of Fully Mixed Bayesian Nash Equilibria

In this section, we study fully mixed Bayesian Nash equilibria for the case of identical links. We start by proving a technical lemma that will be handy later on (Lemma 6.5). With the help of this lemma, we prove a simple expression on the private cost of the players in a fully mixed Bayesian Nash equilibrium (Theorem 6.6). Then, we show that fully mixed Bayesian Nash equilibria maximize private costs (Theorem 6.7). This result will be of particular interest in Section 6.4. We proceed with an exact characterization of fully mixed Bayesian Nash equilibria (Theorem 6.9). Finally, we determine the dimension of space of fully mixed Bayesian Nash equilibria (Theorem 6.10).

We start with the following technical lemma.

Lemma 6.5. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated mixed strategy profile $\mathbf{P}$. Then, for each player $i \in [n]$,*

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

Proof. Clearly,

$$\begin{aligned}
&\sum_{j\in[m]} \delta_j^{-i}(\mathbf{P},(\mathbf{\Psi}|t_i=t)) \\
&= \sum_{j\in[m]}\sum_{\boldsymbol{\sigma}\in\Sigma}\prod_{s\in[n]} p(s,\sigma_s)\cdot\delta_j^{-i}(\boldsymbol{\sigma},(\mathbf{\Psi}|t_i=t)) \\
&= \sum_{j\in[m]}\sum_{\boldsymbol{\sigma}\in\Sigma}\prod_{s\in[n]} p(s,\sigma_s)\cdot\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} \Psi(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t)\sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w(t_s) \\
&= \sum_{\boldsymbol{\sigma}\in\Sigma}\prod_{s\in[n]} p(s,\sigma_s)\cdot\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} \Psi(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t)\sum_{s\in[n]\setminus\{i\}} w(t_s) \\
&= \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} \Psi(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t)\sum_{s\in[n]\setminus\{i\}} w(t_s) \\
&= \sum_{s\in[n]\setminus\{i\}}\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} \Psi(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t)w(t_s) \\
&= \sum_{s\in[n]\setminus\{i\}} W(s|t_i=t),
\end{aligned}$$

as needed. ■

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

Theorem 6.6. *Consider a Bayesian routing game $\Gamma=(n,m,\mathbf{1},T,\mathbf{\Psi})$ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Then for each player $i\in[n]$,*

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

Proof. Fix any player $i\in[n]$. Clearly, for any link $k\in\mathsf{support}_i(\mathbf{F})=[m]$, and by Lemma 6.5,

$$\begin{aligned}
\mathsf{PC}_i(\mathbf{F},\mathbf{\Psi}) &= \sum_{t\in T_i}\Psi(i,t)\cdot v_{(i,t)}(\mathbf{F},\mathbf{\Psi}) \\
&= \sum_{t\in T_i}\Psi(i,t)\cdot\big(w(t)+\delta_k^{-i}(\mathbf{F},(\mathbf{\Psi}|t_i=t))\big) \\
&= \sum_{t\in T_i}\Psi(i,t)\cdot w(t) + \sum_{t\in T_i}\Psi(i,t)\cdot\delta_k^{-i}(\mathbf{F},(\mathbf{\Psi}|t_i=t)) \\
&= W(i) + \sum_{t\in T_i}\Psi(i,t)\cdot\frac{1}{m}\cdot\sum_{j\in[m]}\delta_j^{-i}(\mathbf{F},(\mathbf{\Psi}|t_i=t)) \\
&= W(i) + \frac{1}{m}\sum_{t\in T_i}\Psi(i,t)\sum_{s\in[n]\setminus\{i\}} W(s|t_i=t) \qquad (6.1)
\end{aligned}$$

$$
\begin{aligned}
&= W(i) + \frac{1}{m} \sum_{s \in [n] \setminus \{i\}} \sum_{t \in T_i} \Psi(i,t) \cdot W(s|t_i = t) \\
&= W(i) + \frac{1}{m} \sum_{s \in [n] \setminus \{i\}} W(s) \\
&= \frac{W}{m} + \frac{m-1}{m} W(i),
\end{aligned}
$$

as needed. ■

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 was shown by Gairing et al. [47].

Theorem 6.7. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and Bayesian Nash equilibrium* $\mathbf{P}$. *Then for each player $i \in [n]$,*

$$
\mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi}) \leq \mathsf{PC}_i(\mathbf{F}, \mathbf{\Psi}).
$$

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

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

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

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

by Equation (6.1), as needed. ■

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

Definition 6.8. *The* standard *fully mixed strategy profile* $\mathbf{F}^*$ *is the 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, the standard fully mixed strategy profile is a Bayesian Nash equilibrium. For the special case of complete information routing games, this fact was first stated in [74].

In general, there exists more than one fully mixed Bayesian Nash equilibrium. In the remainder of this section, we study the structure of fully mixed Bayesian Nash equilibria for Bayesian routing games on identical links with independent type distribution. We start with an exact characterization of fully mixed Bayesian Nash equilibria.

Theorem 6.9. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ 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*

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

for all players $i \in [n]$ and links $j \in [m]$.

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

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

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

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

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

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

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

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

So, $v_{(i,t)}(\mathbf{F},\mathbf{\Psi}) = \lambda^j_{(i,t)}(\mathbf{F},\mathbf{\Psi})$ 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.

We will now show the opposite direction. Assume that $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium. It follows that $\mathsf{support}_t(\mathbf{P}) = [m]$ for all players $i\in[n]$ and types $t\in T_i$. Since $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium and $\mathbf{\Psi}$ is independent, it follows that for all links $j\in\mathsf{support}_t(\mathbf{P}) = [m]$,

$$\begin{aligned}v_{(i,t)}(\mathbf{F},\mathbf{\Psi}) &= \lambda^j_{(i,t)}(\mathbf{F},\mathbf{\Psi})\\ &= \delta^{-i}_j(\mathbf{F},(\mathbf{\Psi}|t_i = t)) + w(t)\\ &= \delta^{-i}_j(\mathbf{F},\mathbf{\Psi}) + w(t).\end{aligned}$$

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

$$\delta^{-i}_j(\mathbf{F},\mathbf{\Psi}) = \delta^{-i}_l(\mathbf{F},\mathbf{\Psi}).$$

Since $\delta^{-i}_j(\mathbf{F},\mathbf{\Psi}) = \sum_{k\in[n]\setminus\{i\}}\mu(\mathbf{F},k,j)$ 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 $j_1\neq j_2$,

$$\sum_{k\in[n]\setminus\{i_1\}} \mu(\mathbf{F},k,j_1) = \sum_{k\in[n]\setminus\{i_1\}} \mu(\mathbf{F},k,j_2) \tag{6.2}$$

and

$$\sum_{k\in[n]\setminus\{i_2\}} \mu(\mathbf{F},k,j_1) = \sum_{k\in[n]\setminus\{i_2\}} \mu(\mathbf{F},k,j_2). \tag{6.3}$$

Subtracting (6.3) from (6.2) 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

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

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 $i \in [n]$,

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

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

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

as needed. ■

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

Theorem 6.10. *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 $\sum_{i \in [n]} m^{\tau_i} - nm$.*

Proof. Let **F** be a fully mixed Bayesian Nash equilibrium. By Theorem 6.9, this is equivalent to **F** being a fully mixed strategy profile and

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

for all players $i \in [n]$ and links $j \in [m]$. So, **F** is a solution to the system of linear equations and inequalities:

$$\begin{aligned}
&(1)\ f(i, \sigma_i) > 0 && \forall i \in [n], \forall \sigma_i \in \Sigma_i \\
&(2)\ \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) = 1 && \forall i \in [n] \\
&(3)\ \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \sum_{t \in T_i : \sigma_i(t) = j} \Psi(i, t) \cdot w(t) = \frac{1}{m} \cdot W(i) && \forall i \in [n], \forall j \in [m].
\end{aligned}$$

The dimension of the solution space of this system is the number of variables minus the number of independent equations. For each player $i \in [n]$ we have m^{τ_i} variables. Thus, the total number of variables is $\sum_{i \in [n]} m^{\tau_i}$. We now show an upper bound on the number of independent equations. Fix any player $i \in [n]$. Summing up the equations (3) for all links $j \in [m]$ yields

$$\begin{aligned}
& \sum_{j \in [m]} \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \sum_{t \in T_i : \sigma_i(t) = j} \Psi(i, t) \cdot w(t) = \sum_{j \in [m]} \frac{1}{m} \cdot W(i) \\
\Leftrightarrow \quad & \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \sum_{t \in T_i} \Psi(i, t) \cdot w(t) = W(i) \\
\Leftrightarrow \quad & \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot W(i) = W(i) \\
\Leftrightarrow \quad & \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) = 1
\end{aligned}$$

It follows that all equations (2) are implied by a linear combination of equations in (3). Therefore, nm is an upper bound on the number of independent equations. The claim follows. ∎

6.4 Social Cost and Price of Anarchy

In this section, we present bounds on the price of anarchy for three different social cost measures. All these results are for the case of *identical* links. In Section 6.4.1, we summarize our results for makespan social cost. In Section 6.4.2, we consider social cost as sum of private cost and in Section 6.4.3, we present our findings for social cost as maximum of private costs.

6.4.1 Makespan Social Cost

In this section, we study the price of anarchy for *makespan social cost*. For the special case of complete information routing games this social cost measure was introduced in [67] and asymptotic tight bounds on the price of anarchy were given by Czumaj and Vöcking [23] and Koutsoupias et al. [66]. 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 previous techniques cannot be applied to prove an upper bound on the price of anarchy which is better than $\mathcal{O}(m)$.

Proposition 6.11. *For any $\epsilon > 0$, there is a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links with independent type distribution and an associated pure Bayesian Nash equilibrium $\boldsymbol{\sigma}$ with $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma}) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma)$, such that for each link $j \in [m]$,*

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

Proof. Set $n = m$. For each player $i \in [n]$, set $T_i = \{t_i, t_i'\}$ with $w(t_i) = 0$ and $w(t_i') = a$; set also for each player $i \in [n]$, $\Psi(i, t_i) = 1 - \frac{1}{a}$ and $\Psi(i, t_i') = \frac{1}{a}$. Let $\boldsymbol{\sigma}$ be the pure Bayesian Nash equilibrium that maps both types of player i to link i, where $i \in [n]$. Since each player is assigned to a different link, we have $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma})$. Clearly, on the one hand, $\delta_j(\boldsymbol{\sigma}, \mathbf{\Psi}) = 1$ for all links $j \in [m]$. On the other hand,

$$\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma}) = \left(1 - \left(1 - \frac{1}{a}\right)^m\right) a.$$

Note that

$$\begin{aligned}
\lim_{a\to\infty}\left(\left(1 - \left(1 - \frac{1}{a}\right)^m\right) a\right) &= \lim_{a\to\infty}\left(\left(1 - \sum_{i=0}^{m}\binom{m}{i}\left(-\frac{1}{a}\right)^i\right) a\right) \\
&= \lim_{a\to\infty}\left(\left(1 - 1 - \sum_{i=1}^{m}\binom{m}{i}(-1)^i\left(\frac{1}{a}\right)^i\right) a\right) \\
&= \lim_{a\to\infty}\left(\sum_{i=1}^{m}\binom{m}{i}(-1)^{i-1}\left(\frac{1}{a}\right)^{i-1}\right) \\
&= \lim_{a\to\infty}\left(m + \sum_{i=2}^{m}\binom{m}{i}(-1)^{i-1}\left(\frac{1}{a}\right)^{i-1}\right) \\
&= m.
\end{aligned}$$

The claim follows. ■

We now turn our attention to the standard fully mixed Bayesian Nash equilibrium on identical links. We prove:

Theorem 6.12. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated standard fully mixed Bayesian Nash equilibrium $\mathbf{F}^*$. Then,*

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

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 $\overline{\mathbf{P}}(t)$ assigns each player to each link with probability $\frac{1}{m}$ (see [74, Lemma 15]). By [66, Theorem 4.4] or [23, Theorem 1.1], it holds that

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

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

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

as needed. ■

Theorem 6.12 implies that for the standard fully mixed Nash equilibrium, incomplete information has no impact on the price of anarchy if social cost is taken as makespan social cost.

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

Proposition 6.13. *There exists a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$ such that*

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

Proof. Consider the Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ 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 $\Psi(i, t_i) = \Psi(i, t_i') = \frac{1}{2}$ for all players $i \in \{1, 2\}$. Consider the standard fully mixed Bayesian Nash equilibrium $\mathbf{F}^*$ and some other fully mixed Bayesian Nash equilibrium $\mathbf{F}$ which we define below:

- $\mathbf{F}^*$ 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 $\mathbf{F}^*$ is $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \mathbf{F}^*) = \frac{1}{3} \cdot \left(\frac{1}{4} \cdot 2 + \frac{1}{2} \cdot 3 + \frac{1}{4} \cdot 4\right) + \frac{2}{3} \cdot \left(\frac{1}{4} \cdot 1 + \frac{3}{4} \cdot 2\right) = \frac{13}{6}$.
- The fully mixed strategy profile $\mathbf{F}$ assigns each type of traffic 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 traffic 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 traffic 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 traffic 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 traffic 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 $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \mathbf{F}) = \frac{1}{4} \cdot \left(\frac{3}{8} \cdot 2 + \frac{5}{8} \cdot 1\right) + \frac{1}{2} \cdot \left(\frac{5}{16} \cdot 3 + \frac{11}{16} \cdot 2\right) + \frac{1}{4} \cdot \left(\frac{11}{32} \cdot 4 + \frac{21}{32} \cdot 2\right)$ $= \frac{139}{64}$.

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

It is known (see [73, 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 6.14. *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 traffic 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 [23, 66] and related links [23]. We use these bounds to prove:

Theorem 6.15. *Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{\Psi})$ be a CI-like game with an associated pure Bayesian Nash equilibrium $\boldsymbol{\sigma}$. Then*

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

(b) $\frac{\mathsf{SC}_{\mathsf{MSP}}(\Gamma,\boldsymbol{\sigma})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log \log \log m}\right)$*, 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 $\boldsymbol{\sigma}$ to a complete information routing Γ_{CI} with associated (mixed) strategy profile $\mathbf{P}$. For this construction, we show that $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma}) = \mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}, \mathbf{P})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$, and that $\mathbf{P}$ is a Nash equilibrium if $\boldsymbol{\sigma}$ is a Bayesian Nash equilibrium. From these properties of our construction, we derive that the corresponding upper bounds on the price of anarchy [23, 66] 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 $\mathbf{P}$, we can define a CI-like game Γ with associated pure Bayesian Nash equilibrium $\boldsymbol{\sigma}$, such that our construction maps Γ and $\boldsymbol{\sigma}$ to Γ_{CI} and $\mathbf{P}$, 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 $\Gamma = (n, m, \mathbf{c}, T, \boldsymbol{\Psi})$ be a CI-like game. For each $i \in [n]$, denote by $w_i = w(t)$ the traffic 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 $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ 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 $\mathbf{P}$ for Γ_{CI}, where for each player $i \in [n]$ and all pure strategies $\sigma'_i \in \Sigma'_i$ the probability $p(i, \sigma'_i)$ is given by $p(i, \sigma'_i) = \sum_{t \in T_i : \sigma_i(t) = \sigma'_i(t'_i)} \Psi(i, t)$.

We proceed by showing properties of our construction.

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

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

$$= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\prod_{i\in[n]} \sum_{\substack{t\in T_i:\\ \sigma_i(t)=\sigma'_i(t'_i)}} \Psi(i,t) \right) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i \right\}$$

$$= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \prod_{i\in[n]} p(i,\sigma'_i) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i \right\}$$

$$= \mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}, \mathbf{P}).$$

- $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$: 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 $\hat{\boldsymbol{\sigma}}'$ for Γ_{CI} of optimum social cost, i.e. $\mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}, \hat{\boldsymbol{\sigma}}') = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$. Consider the normal pure strategy profile $\hat{\boldsymbol{\sigma}}$ for Γ that assigns for each $i \in [n]$ all types of player i to the link to that $\hat{\boldsymbol{\sigma}}'$ 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 $\hat{\boldsymbol{\sigma}}$ back to Γ_{CI} and $\hat{\boldsymbol{\sigma}}'$. Thus, $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \hat{\boldsymbol{\sigma}}) = \mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}, \hat{\boldsymbol{\sigma}}')$. We get that

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

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

$$\lambda^j_{(i,t)}(\boldsymbol{\sigma}, \boldsymbol{\Psi})$$

$$= \frac{1}{c_j} \cdot \left(w(t) + \delta_j^{-i}(\boldsymbol{\sigma}, \boldsymbol{\Psi}) \right)$$

$$= \frac{1}{c_j} \cdot \left(w(t) + \sum_{(t_1,\ldots,t_n)\in T} \Psi(t_1,\ldots,t_n) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w(t_s) \right)$$

$$= \frac{1}{c_j} \cdot \left(w_i + \sum_{(t_1,\ldots,t_n)\in T} \prod_{s\in[n]} \Psi(s,t_s) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w_s \right)$$

$$= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ \sigma_s(t_s)=\sigma'_s(t'_s)\forall s\in[n]}} \prod_{s\in[n]} \Psi(s,t_s) \right) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w_s \right)$$

$$= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \left(\prod_{s\in[n]} \sum_{\substack{t_s\in T_s:\\ \sigma_s(t_s)=\sigma'_s(t'_s)}} \Psi(s,t_s) \right) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w_s \right)$$

$$= \frac{1}{c_j} \cdot \left(w(t'_i) + \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'} \prod_{s\in[n]} p(s,\sigma'_s) \cdot \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma'_s(t'_s)=j}} w(t'_s) \right)$$

$$= \frac{1}{c_j} \cdot \left(w(t'_i) + \delta_j^{-i}(\mathbf{P},1) \right)$$

$$= \lambda^j_{(i,t'_i)}(\mathbf{P},1).$$

We now use this property to show that $\mathbf{P}$ is a Nash equilibrium for Γ_{CI} if $\boldsymbol{\sigma}$ is a pure Bayesian Nash equilibrium for Γ. So, let $\boldsymbol{\sigma}$ 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 traffic. Thus,

$$v_{(i,t)}(\boldsymbol{\sigma},\boldsymbol{\Psi}) = v_{(i,\hat{t})}(\boldsymbol{\sigma},\boldsymbol{\Psi})$$

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

$$\begin{aligned} v_{(i,t)}(\boldsymbol{\sigma},\boldsymbol{\Psi}) &= \lambda^j_{(i,t)}(\boldsymbol{\sigma},\boldsymbol{\Psi}) \text{ for all } j \in \mathsf{support}_i(\boldsymbol{\sigma}) \text{ and} \\ v_{(i,t)}(\boldsymbol{\sigma},\boldsymbol{\Psi}) &\le \lambda^j_{(i,t)}(\boldsymbol{\sigma},\boldsymbol{\Psi}) \text{ for all } j \notin \mathsf{support}_i(\boldsymbol{\sigma}). \end{aligned}$$

By definition of $\mathbf{P}$,

$$\mathsf{support}_i(\boldsymbol{\sigma}) = \mathsf{support}_{t'_i}(\mathbf{P}).$$

It follows that

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

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

Upper bounds on price of anarchy: Recall that by our construction, we have that $\mathsf{SC}_{\mathsf{MSP}}(\Gamma,\boldsymbol{\sigma}) = \mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}},\mathbf{P})$ 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 [66] and [23], we get

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

This completes the proof of the upper bounds.

Tightness of the upper bounds: From [66] and [23], there exist complete information routing games Γ_{CI} with an associated mixed Nash equilibrium $\mathbf{P}$ such that

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

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 $\mathbf{P}$. With a slight abuse of notation, we denote $\mathbf{P} = (p(i,j))_{i\in[n],j\in[m]}$ where $p(i,j)$ is the probability that type $t'_i \in T'_i$ is assigned to link $j \in [m]$.

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

> For each player $i \in [n]$, T_i consists of $|\mathsf{support}_i(\mathbf{P})|$ types, where we have a type t_i^j for every link $j \in \mathsf{support}_i(\mathbf{P})$. For all players $i \in [n]$ and links $j \in \mathsf{support}_i(\mathbf{P})$, define $\Psi(i, t_i^j) = p(i,j)$ 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 $\boldsymbol{\sigma}$ back to the complete information routing game Γ_{CI} with associated (mixed) Nash equilibrium $\mathbf{P}$. It follows that $\mathsf{SC}_{\mathsf{MSP}}(\Gamma,\boldsymbol{\sigma}) = \mathsf{SC}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}},\mathbf{P})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$ and $\lambda^l_{(i,t_i^j)}(\boldsymbol{\sigma},\mathbf{\Psi}) = \lambda^l_{(i,t'_i)}(\mathbf{P},1)$ for all players $i \in [n]$, for all links $l \in [m]$, and for all $j \in \mathsf{support}_i(\mathbf{P})$. Since $\mathbf{P}$ is a Nash equilibrium we have

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

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

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

This completes the proof. ■

We conclude with a lower bound on the price of anarchy for social cost as expected maximum latency if we restrict to normal pure Bayesian Nash equilibria.

Theorem 6.16. *There exists a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated normal pure Bayesian Nash equilibrium $\boldsymbol{\sigma}$ such that*

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

Proof. Let $m \in \mathbb{N}$ be a perfect square. Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links with independent type distribution $\mathbf{\Psi}$. There are 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 $w(t_i) = 1, w(t'_i) = 0, \Psi(i, t_i) = \frac{1}{\sqrt{m}}$ and $\Psi(i, t'_i) = 1 - \frac{1}{\sqrt{m}}$ for all players $i \in \mathcal{U}_1$.
- The class $\mathcal{U}_2$ consists of $(\sqrt{m} - 1)m$ players with type set $T_i = \{t_i\}$, where $w(t_i) = \frac{1}{\sqrt{m}}$ and $\Psi(i, t_i) = 1$ all players $i \in \mathcal{U}_2$.

Consider the pure strategy profile $\boldsymbol{\sigma}'$ 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 $\boldsymbol{\sigma}'$, we get

$$\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma}') \leq 1 + (\sqrt{m} - 1) \cdot \frac{1}{\sqrt{m}} < 2.$$

Now consider the normal pure strategy profile $\boldsymbol{\sigma}$ 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, $\boldsymbol{\sigma}$ is a normal pure Bayesian Nash equilibrium.

To show a lower bound on $\mathsf{SC}_{\mathsf{MSP}}(\Gamma, \boldsymbol{\sigma})$ 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 $\mathbf{E}(X_j) = 1$. Let $1 \leq k \leq \sqrt{m}, k \in \mathbb{N}$; the precise choice of k will be made later. Clearly,

$$\begin{aligned}
\mathbf{Pr}(X_j \geq k) &\geq \mathbf{Pr}(X_j = k) \\
&= \binom{\sqrt{m}}{k} \cdot \left(\frac{1}{\sqrt{m}}\right)^k \cdot \left(1 - \frac{1}{\sqrt{m}}\right)^{\sqrt{m}-k} \\
&\geq \binom{\sqrt{m}}{k} \cdot \left(\frac{1}{\sqrt{m}}\right)^k \cdot \frac{1}{\mathsf{e}} \qquad\qquad \text{(since } k \geq 1\text{)} \\
&= \frac{\sqrt{m} \cdot \ldots \cdot (\sqrt{m} - k + 1)}{\sqrt{m}^k} \cdot \frac{1}{k!} \cdot \frac{1}{\mathsf{e}}.
\end{aligned}$$

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

$$\frac{\sqrt{m} \cdot \ldots \cdot (\sqrt{m} - k + 1)}{\sqrt{m}^k} \geq \frac{k!}{k^k}.$$

It follows that

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

so that

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

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

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

Define now $\alpha > 0$ so that $\left(\frac{\alpha}{\mathsf{e}}\right)^{\alpha} = m$. Then, clearly, $\alpha = \Theta\left(\frac{\log m}{\log \log m}\right)$. Choose $k = \frac{\alpha}{\mathsf{e}}$. Then $k^k = m^{\frac{1}{\mathsf{e}}}$ which implies

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

for suitably large m. This implies that

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

Thus,

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

as needed. ■

6.4.2 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.

Theorem 6.7 implies that fully mixed Bayesian Nash equilibria maximize social cost as sum of private costs. Hence, we obtain:

Theorem 6.17. *Consider a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and a Bayesian Nash equilibrium* $\mathbf{P}$*. Then,*

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

We now use Theorem 6.17 to prove an asymptotically tight bound on the price of anarchy for the case of identical links (and social cost as sum of private costs).

Theorem 6.18. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated Bayesian Nash equilibrium* $\mathbf{P}$*. Then,*

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

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 6.17, it suffices to prove the upper bound for a fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Clearly, on the one hand,

$$\begin{aligned}
\mathsf{SC}_{\mathsf{SUM}}(\Gamma, \mathbf{F}) &= \sum_{i \in [n]} \mathsf{PC}_i(\mathbf{F}, \mathbf{\Psi}) \\
&= \sum_{i \in [n]} \left(\frac{W}{m} + \frac{m-1}{m} W(i) \right) \quad \text{(by Proposition 6.6)} \\
&= \frac{nW}{m} + \frac{m-1}{m} W \\
&= \frac{m+n-1}{m} W.
\end{aligned}$$

On the other hand, $\mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi}) \geq W(i)$ for any player $i \in [s]$ and any strategy profile $\mathbf{P}$; 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 traffic, we get

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

Let $\boldsymbol{\sigma}$ 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,

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

Clearly, for any $\varepsilon > 0$, there is a $k \in \mathbb{N}$ such that $\frac{(n-m)(n-m+1)}{n+(m-1)(k-1)} \leq \varepsilon$. 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 $\mathsf{SC}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}, \mathbf{F}) = \frac{m+n-1}{m} W$, it follows that

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

as needed. ■

Berenbrink et al. [11] showed that the price of anarchy for complete information routing games 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 6.18 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 6.18 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 of private costs, as we will see next.

6.4.3 Social Cost as Maximum of Private Costs

In this section, we study the price of anarchy for social cost as the maximum of private costs.

Theorem 6.7 implies that fully mixed Bayesian Nash equilibria maximize social cost as maximum of private costs. Hence, we obtain:

Theorem 6.19. *Consider a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and a Bayesian Nash equilibrium* $\mathbf{P}$*. Then,*

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

We now use Theorem 6.19 to prove asymptotically tight bounds on the price of anarchy for the case of identical links.

Theorem 6.20. *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ on identical links and an associated Bayesian Nash equilibrium $\mathbf{P}$. Then,*

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

(b) $\dfrac{\mathsf{SC}_{\mathsf{MAX}}(\Gamma, \mathbf{P})}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \leq 2 - \dfrac{1}{m}$, *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. Theorem 6.6 and Theorem 6.7 imply together that

$$\mathsf{PC}_i(\mathbf{P}, \mathbf{\Psi}) \leq \mathsf{PC}_i(\mathbf{F}, \mathbf{\Psi}) = \frac{W}{m} + \frac{m-1}{m} W(i), \tag{6.4}$$

for each player $i \in [n]$. We consider the two cases from the theorem.
Case (a):
Upper bound: Clearly, for any strategy profile $\mathbf{P}'$ and for any player $i \in [n]$, $\mathsf{PC}_i(\mathbf{P}', \mathbf{\Psi}) \geq W(i)$; hence, $\sum_{i \in [n]} \mathsf{PC}_i(\mathbf{P}', \mathbf{\Psi}) \geq W$. This implies that

$$\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \geq \frac{W}{n}. \tag{6.5}$$

Clearly, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \geq W(i)$ for all $i \in [n]$. Fix any player $i \in [n]$. By (6.4) and (6.5),

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

and the upper bound follows.
Lower bound: Fix any arbitrary $k, a, r \in \mathbb{N}$, which will be determined later. Consider the Bayesian routing game $\Gamma_{k,a,r} = (n, m, \mathbf{1}, T, \mathbf{\Psi})$ 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 traffic $w(t_i) = 1$, $w(t'_i) = a \cdot r$ and probabilities $\Psi(i, t_i) = 1 - \frac{1}{a}$, $\Psi(i, t'_i) = \frac{1}{a}$. Clearly, for player $i \in [n]$, $W(i) = r + 1 - \frac{1}{a}$.

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

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

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

$$\mathsf{PC}_i(\boldsymbol{\sigma}, \boldsymbol{\Psi}) \leq (k+r)(1+\epsilon').$$

Hence, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r}) \leq (k+r)(1+\epsilon')$. Fix now any fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Theorem 6.6 implies that for each player $i \in [n]$,

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

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

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

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

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

This proves that the upper bound shown before is tight up to a factor of $(1+\epsilon)$.
Case (b):
Upper bound:
Consider the 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 (6.4),

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

so that

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

as needed. The upper bound follows.
Lower bound:
Consider 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 (6.4),

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

so that

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

as needed. ∎

6.5 Conclusion and Discussion

In this chapter, we have introduced the dimension of incomplete information into the class of routing games on parallel links. For this setting, we have studied the existence and computational complexity of pure Bayesian Nash equilibria, structural properties of fully mixed Bayesian Nash equilibria and the price of anarchy for different social cost measures.

Our work leaves open several interesting problems. On the most concrete level, we would like to ask:

- Can pure Bayesian Nash equilibria be computed in polynomial time?
- What is the exact value of the price of anarchy for identical links if social cost is defined as expected maximum latency?
- What is the price of anarchy for all three considered social cost measures in the case of related links?

References

[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 (FOCS'06)*, 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 (WINE'06)*, 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 B. Durand and W. Thomas, editors, *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS'06)*, Lecture Notes in Computer Science, Vol. 3884, Springer Verlag, pages 218–229, 2006.

[5] 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 (STOC'05)*, pages 57–66, 2005.

[6] B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. In K. Jansen and R. Solis-Oba, editors, *Proceedings of the 1st International Workshop on Approximation and Online Algorithms (WAOA'03)*, Lecture Notes in Computer Science, Vol. 2909, Springer Verlag, pages 41–52, 2003.

[7] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[8] M. J. Beckmann. On the Theory of Traffic Flow in Networks. *Traffic Quarterly*, 21:109–116, 1967.

[9] R. Beier, A. Czumaj, P. Krysta, and B. Vöcking. Computing Equilibria for Congestion Games with (Im)perfect Information. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'04)*, pages 746–755, 2004.

[10] E. T. Bell. Exponential Polynomials. *Annals of Mathematics*, 35:258–277, 1934.

[11] P. Berenbrink, L. A. Goldberg, P. W. Goldberg, and R. Martin. Utilitarian Resource Assignment. *Journal of Discrete Algorithms*, (http://dx.doi.org/10.1016/j.jda.2005.06.009), to appear.

[12] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Second Edition. Springer Verlag, 1980.

[13] D. Braess. Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.

[14] P. Brucker, J. Hurink, and F. Werner. Improving Local Search Heuristics for Some Scheduling Problems. Part II. *Discrete Applied Mathematics*, 72 (1–2):47–69, 1997.

[15] X. Chen and X. Deng. Settling the Complexity of 2-Player Nash-Equilibrium. In *Electronic Colloquium in Computational Complexity*, TR05-140, 2005.

[16] G. Christodoulou and E. Koutsoupias. On The Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In S. Leonardi G. S. Brodal, editor, *Proceedings of the 13th Annual European Symposium on Algorithms (ESA'05)*, Lecture Notes in Computer Science, Vol. 3669 Springer Verlag, pages 59–70, 2005.

[17] 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 (STOC'05)*, pages 67–73, 2005.

[18] 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 (SODA'06)*, pages 668–677, 2006.

[19] R. Comminetti, J. R. Correa, and N. E. Stier-Moses. Network Games with Atomic Players. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP'06), Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 525–536, 2006.

[20] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[21] A. Czumaj. Selfish Routing on the Internet. In J. Leung, editor, *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*. CRC Press, 2004.

[22] 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 (STOC'02)*, pages 287–296, 2002.

[23] A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, pages 413–420, 2002. Also accepted to *Journal of Algorithms* as Special Issue of SODA'02.

[24] D. Dumrauf and M. Gairing. Price of Anarchy for Polynomial Wardrop Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE'06)*, 2006.

[25] R. Elsässer, M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. A Simple Graph-Theoretic Model for Selfish Restricted Scheduling. In X. Deng and Y. Ye, editors, *Proceedings of the 1st International Workshop on Internet and Network Economics (WINE'05)*, Lecture Notes in Computer Science, Vol. 3828, Springer Verlag, pages 195–209, 2005.

[26] E. Even-Dar, A. Kesselmann, and Y. Mansour. Convergence Time to Nash Equilibria. In J. C. M. Baeten, J. K. Lenstra, J. Parrow, and G. J. Woeginger, editors, *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, Lecture Notes in Computer Science, Vol. 2719, Springer Verlag, pages 502–513, 2003.

[27] 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 (STOC'04)*, pages 604–612, 2004.

[28] G. Facchini, F. van Megan, P. Borm, and S. Tijs. Congestion Models and Weighted Bayesian Potential Games. *Theory and Decision*, 42:193–206, 1997.

[29] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. *Journal of Computer and System Sciences*, 63(1): 21–41, 2001.

[30] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the Coordination Ratio for a Selfish Routing Game. In J. C. M. Baeten, J. K. Lenstra, J. Parrow, and G. J. Woeginger, editors, *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, Lecture Notes in Computer Science, Vol. 2719, Springer Verlag, pages 514–526, 2003.

[31] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish Routing in Non-Cooperative Networks: A Survey. In B. Rovan and P. Vojtás, editors, *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 21–45, 2003.

[32] G. Finn and E. Horowitz. A Linear Time Approximation Algorithm for Multiprocessor Scheduling. *BIT*, 19(3):312–320, 1979.

[33] 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 (STOC'06)*, pages 653–662, 2006.

[34] S. Fischer and B. Vöcking. On the Evolution of Selfish Routing. In T. Radzik S. Albers, editor, *Proceedings of the 13th Annual European Symposium on Algorithms (ESA'04)*, Lecture Notes in Computer Science, Vol. 3221 Springer Verlag, pages 323–334, 2004.

[35] S. Fischer and B. Vöcking. Adaptive Routing with Stale Information. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC'05)*, pages 276–283, 2005.

[36] S. Fischer and B. Vöcking. On the Structure and Complexity of Worst-Case Equilibria. In X. Deng and Y. Ye, editors, *Proceedings of the 1st International*

Workshop on Internet and Network Economics (WINE'05), Lecture Notes in Computer Science, Vol. 3828, Springer Verlag, pages 151–160, 2005.

[37] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In P. Widmayer, F. T. Ruiz, R. M. Bueno, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP'02)*, Lecture Notes in Computer Science, Vol. 2380, Springer Verlag, pages 123–134, 2002.

[38] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 593–605, 2004.

[39] D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms (WAOA'05)*, 2005.

[40] D. K. Friesen. Tighter Bounds for LPT Scheduling on Uniform Processors. *SIAM Journal on Computing*, 16(3):554–560, 1987.

[41] 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 (STOC'04)*, pages 613–622, 2004.

[42] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Polynomial Social Cost. In J. Fiala, V. Koubek, and J. Kratochvíl, editors, *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS'04)*, Lecture Notes in Computer Science, Vol. 3153, Springer Verlag, pages 574–585, 2004.

[43] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Polynomial Social Cost. *Theoretical Computer Science*, 369: 116–135, 2006.

[44] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters*, 16(1): 117–131, 2006.

[45] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 645–657, 2004.

[46] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme Nash Equilibria. In C. Blundo and C. Laneve, editors, *Proceedings of the 8th Italian Conference on Theoretical Computer Science (ICTCS'03)*,

Lecture Notes in Computer Science, Vol. 2841, Springer Verlag, pages 1–20, 2003.

[47] 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.

[48] M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. In L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi, and M. Yung, editors, *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming (ICALP'05)*, Lecture Notes in Computer Science, Vol. 3580, Springer Verlag, pages 51–65, 2005.

[49] 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 (SPAA'05)*, pages 203–212, 2005. Also accepted to *Theory of Computing Systems.*

[50] M. Gairing, B. Monien, and K. Tiemann. Routing (Un-) Splittable Flow in Games with Player-Specific Linear Latency Functions. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP'06), Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 501–512, 2006.

[51] M. Gairing, B. Monien, and A. Woclaw. A Faster Combinatorial Approximation Algorithm for Scheduling Unrelated Parallel Machines. In L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi, and M. Yung, editors, *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming (ICALP'05)*, Lecture Notes in Computer Science, Vol. 3580, Springer Verlag, pages 828–839, 2005. Also accepted to *Theoretical Computer Science.*

[52] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, 1979.

[53] C. Georgiou, T. Pavlides, and A. Philippou. Network Uncertainty in Selfish Routing. In *Proceedings of the 20th IEEE International Parallel & Distributed Processing Symposium (IPDPS'06), CD-ROM Proceedings*, page 105, 2006.

[54] P. W. Goldberg. Bounds on the Convergnece Rate of Randomized Local Search in a Multiplayer Load-balancing Game. In *Proceedings of the 23th Annual ACM Symposium on Principles of Distributed Computing (PODC'04)*, pages 131–140, 2004.

[55] G. H. Gonnet. Expected Length of the Longest Probe Sequence in Hash Code Searching. *Journal of the ACM*, 28(2):289–304, 1981.

[56] R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

[57] R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete Mathematics: A Foundation for Computer Science.* Addison-Wesley, 1994.

[58] J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.

[59] J. C. Harsanyi. Games with Randomly Disturbed Payoffs. *Internation Journal on Game Theory*, 21:1–23, 1973.

[60] J. Havil. *Gamma: Exploring Euler's Constant.* Princeton University Press, 2003.

[61] D. S. Hochbaum and D. B Shmoys. Using Dual Approximation Algorithms for Scheduling Problems: Theoretical and Practical Results. *Journal of the ACM*, 34(1):144–162, 1987.

[62] D. S. Hochbaum and D. B. Shmoys. A Polynomial Approximation Scheme for Scheduling on Uniform Processors: Using the Dual Approximation Approach. *SIAM Journal on Computing*, 17(3):539–551, 1988.

[63] K. Jain and V. Vazirani. Applications of Approximation Algorithms to Cooperative Games. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 364–372, 2001.

[64] R. M. Karp. Reducibility Among Combinatorial Problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, Plenum Press, pages 85–103, 1972.

[65] E. Koutsoupias. Selfish Task Allocation. *Bulletin of the EATCS*, 81:79–88, 2003.

[66] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.

[67] E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In C. Meinel and S. Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, Lecture Notes in Computer Science, Vol. 1563, Springer Verlag, pages 404–413, 1999.

[68] L. Libman and A. Orda. The Designer's Perspective to Atomic Noncooperative Networks. *IEEE/ACM Transactions on Networking*, 7(6):875–884, 1999.

[69] L. Libman and A. Orda. Atomic Resource Sharing in Noncooperative Networks. *Telecommunication Systems*, 17(4):385–409, 2001.

[70] T. Lücking. Analysing Models for Scheduling and Routing, PhD-thesis, University of Paderborn, 2005.

[71] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In V. Diekert and M. Habib, editors, *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS'04)*, Lecture Notes in Computer Science, Vol. 2996, Springer Verlag, pages 547–558, 2004.

[72] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash Equilibrium? In B. Rovan and P. Vojtás, editors, *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 551–561, 2003.

[73] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[74] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 510–519, 2001.

[75] V. Mazalov, B. Monien, F. Schoppmann, and K. Tiemann. Wardrop Equilibria and Prize of Stability for Bottleneck Games with Splittable Traffic. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE'06)*, 2006.

[76] R. D. McKelvey and A. McLennan. Computation of Equilibria in Finite Games. In H. Amman, D. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, 1996.

[77] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

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

[79] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.

[80] 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.

[81] J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.

[82] N. Nisan. Algorithms for Selfish Agents. In C. Meinel and S. Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, Lecture Notes in Computer Science, Vol. 1563, Springer Verlag, pages 1–15, 1999.

[83] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC'99)*, pages 129–140, 1999.

[84] A. Orda, R. Rom, and N. Shimkin. Competitive Routing in Multiuser Communication Networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

[85] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[86] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 749–753, 2001.

[87] A. C. Pigou. *The Economics of Welfare*. Macmillan and Company, 1920.

[88] R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[89] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[90] T. Roughgarden. *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[91] T. Roughgarden. Selfish Routing with Atomic Players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 1184–1185, 2005.

[92] T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.

[93] P. Schuurman and T. Vredeveld. Performance Guarantees of Local Search for Multiprocessor Scheduling. In K. Aardal and B. Gerards, editors, *Proceedings of the 8th Conference on Integer Programming and Combinatorial Optimization (IPCO'01)*, Lecture Notes in Computer Science, Vol. 2081, Springer Verlag, pages 370–382, 2001.

[94] J. Stirling. Methodus Differentialis Sive Tractatus de Summatione et Interpolatione Serierum Infinitarum. *London*, 1730.

[95] S. Suri, C. D. Tóth, and Y. Zhou. Selfish Load Balancing and Atomic Congestion Games. In *Proceedings of the 16th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'04)*, pages 188–195, 2004.

[96] 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.

[97] J. G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proceedings of the Institute of Civil Engineers, Pt. II, Vol. 1*, pages 325–378, 1952.