

Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture ^{*}

Martin Gairing, Thomas Lücking, Burkhard Monien, and Karsten Tiemann^{**}

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{gairing,luck,bm,tiemann}@uni-paderborn.de

1 Introduction

Motivation-Framework. Apparently, it is in human’s nature to act selfishly. Game Theory, founded by von Neumann and Morgenstern [39, 40], provides us with *strategic games*, an important mathematical model to describe and analyze such a selfish behavior and its resulting conflicts. In a strategic game, each of a finite set of players aims for an optimal value of its *private objective function* by choosing either a *pure* strategy (a single strategy) or a *mixed* strategy (a probability distribution over all pure strategies) from its *strategy set*. Strategic games in which the strategy sets are finite are called *finite strategic games*. Each player chooses its strategy once and for all, and all players’ choices are made *non-cooperatively* and *simultaneously* (that is, when choosing a strategy each player is not informed of the strategies chosen by any other player). One of the basic assumption in strategic games is that the players act *rational*, that is, consistently in pursuit of their private objective function. For a concise introduction to contemporary Game Theory we recommend [25].

One of the most widely used solution concepts for strategic games is the concept of *Nash equilibrium*. It represents a stable state in which no player wishes to leave unilaterally its own strategy in order to improve the value of its private objective function. A Nash equilibrium is called *pure* if all players choose a pure strategy, otherwise *mixed*. Many algorithms have been developed to compute a Nash equilibrium (see [27] for an overview). Though the celebrated results of Nash [30, 31] ensure the existence of a mixed Nash equilibrium, the complexity to compute such a Nash equilibrium is widely unknown. Papadimitriou [32] advocates it to be “*the most important concrete open question on the boundary of \mathcal{P} today*”.

Rosenthal [33] introduced a special class of strategic games, now widely known as *congestion games*. Here, the strategy set of each player is a subset of the power set of given *resources*. The players share a private objective function, defined as the sum (over their chosen resources) of functions in the *number* of players sharing this resource. In his seminal work, Rosenthal [33] showed with help of a *potential function* that congestion games (in sharp contrast to general strategic games) always admit at least one pure Nash equilibrium. Later, Milchtaich [28] considered two extensions of congestion

^{*} This work has been partially supported by the DFG-SFB 376 and by the European Union within the 6th Framework Programme under contract 001907 (DELIS).

^{**} International Graduate School of Dynamic Intelligent Systems

games, namely *weighted congestion games* in which the players have *weights* and thus different influence on the congestion of the resources, and congestion games with *player-specific payoff-functions* in which the players do not share a private objective function.

Another class of (weighted) congestion games are (weighted) *network congestion games* [8, 11] in which the strategy sets correspond to paths in a network. Koutsoupias and Papadimitriou [21] considered a very simple member of this class, now known as *KP-model*. The network consists of a single *source* and a single *destination* which are connected by parallel *links*. Associated with each link is a *capacity* representing the rate at which the link processes *load*, that is, the total weight of players assigned to this link. Thus, the latency functions are linear. Each of the players selfishly routes from the source to the destination by choosing a probability distribution over the links. The private objective function of a player is defined as its expected latency.

Koutsoupias and Papadimitriou [21] were not only interested in the computational complexity of Nash equilibria but also in the degradation of the social welfare of the system due to the selfish behavior of the players. In order to measure this social welfare, they introduced a global objective function, usually coined as *social cost*, which is defined as the expected maximum latency on a link, where the expectation is taken over all random choices of the players. The *price of anarchy*, also called *coordination ratio*, measures the extent to which non-cooperation approximates cooperation. It is defined as the worst-case ratio between the value of social cost in a Nash equilibrium and that of some social optimum. So, the price of anarchy represents a rendezvous of Nash equilibrium, a concept fundamental to Game Theory, with approximation, an ubiquitous concept in Theoretical Computer Science today (see, e.g., [38]).

Mavronicolas and Spirakis [26] introduced the notion of a *fully mixed* Nash equilibrium in which each player chooses every link with positive probability. Gairing *et al.* [15] conjectured that, in case of its existence, the *fully mixed* Nash equilibrium is the worst Nash equilibrium with respect to social cost. This so-called *Fully Mixed Nash Equilibrium Conjecture* is simultaneously intuitive and natural. To support intuition, observe that the fully mixed Nash equilibrium favors collisions between different players (since each player assigns its item with positive probability to every link). This increased probability of collisions should favor an increase to social cost. To support significance, note that the Fully Mixed Nash Equilibrium Conjecture identifies the worst-case Nash equilibrium of *all* instances. We stress that, in sharp contrast, the price of anarchy only determines the worst-case Nash equilibrium of *worst-case* instances.

Recently, the KP-model was extended to *restricted strategy sets* [2, 13] where the strategy set of each player is a *subset* of the links. In addition, the KP-model was extended to general latency functions and studied with respect to different definitions of social cost [1, 14]. Inspired by the arisen interest in the price of anarchy, the much older *Wardrop-model* [3, 6, 41] was re-investigated [35, 36]. In this weighted network congestion game, weight can be split into arbitrary pieces. The social welfare of the system is defined as the sum of the edge latencies. An equilibrium in the Wardrop-model can be interpreted as a Nash equilibrium in a game with infinitely many players, each carrying an infinitesimal amount of weight. Finally, the price of anarchy found its way into congestion games [4, 11].

In this paper, we give a thorough survey on the most exciting results on finite (weighted) congestion games and the special classes mentioned above. In particular, we review the findings on the existence and computational complexity of pure Nash equilibria. Furthermore, we discuss results on the price of anarchy. Last but not least, we survey known facts on fully mixed Nash equilibria.

Overview. The rest of this paper is organized as follows. After a formal definition of (weighted) congestion games in Section 2, we turn our attention to the existence and computational complexity of pure Nash equilibria in Section 3. In Section 4, we consider the price of anarchy before we investigate fully mixed Nash equilibria in Section 5. We conclude, in Section 6, with some open problems.

2 Definitions and Notations

For all integers $k \geq 0$, we denote $[k] = \{1, \dots, k\}$.

A *weighted congestion game* Γ is a tuple

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

Here, n is the number of *players* and E is the finite set of *resources*. For every player $i \in [n]$, w_i is the *weight* and $S_i \subseteq 2^E$ is the *strategy set* of player i . Denote $W = \sum_{i \in [n]} w_i$ and $S = S_1 \times \dots \times S_n$. For every resource $e \in E$, the *latency function* $f_e : \mathbb{R}^+ \mapsto \mathbb{R}^+$ describes the *latency* on resource e .

In a *congestion game*, the weights of all players are equal. Thus, the private cost of a player only depends on the *number* of players choosing the same resources. A congestion game is *symmetric* if the players share a strategy set.

2.1 Strategies and Assignments

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 assignment* is an n -tuple $\mathbf{L} = (s_1, \dots, s_n) \in S$ whereas a *mixed assignment* $\mathbf{P} = (P_1, \dots, P_n)$ is represented by an n -tuple of mixed strategies. A mixed assignment is *fully mixed* if $p(i, s_i) > 0$ for all $i \in [n]$ and $s_i \in S_i$.

2.2 Private Cost

Fix any pure assignment \mathbf{L} , and denote by $l_e(\mathbf{L}) = \sum_{i \in [n], s_i \ni e} w_i$ the *load* on resource $e \in E$. The *private cost* of player $i \in [n]$ is defined by

$$\text{PC}_i(\mathbf{L}) = \sum_{e \in s_i} f_e(l_e(\mathbf{L})).$$

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

$$\text{PC}_i(\mathbf{P}) = \sum_{\mathbf{L} \in S} p(\mathbf{L}) \cdot \text{PC}_i(\mathbf{L}).$$

2.3 Social Cost

Associated with a weighted congestion game Γ and a mixed assignment \mathbf{P} is the *social cost* as a measure of social welfare. We consider the following three definitions of social cost:

- Sum of private costs $SC_{\text{SUM}}(\mathbf{P}) = \sum_{i \in [n]} PC_i(\mathbf{P})$
- Maximum of private costs $SC_{\text{MAX}}(\mathbf{P}) = \max_{i \in [n]} PC_i(\mathbf{P})$
- Expected maximum latency $SC_{\infty}(\mathbf{P}) = \sum_{\mathbf{L} \in S} p(\mathbf{L}) \cdot \max_{i \in [n]} PC_i(\mathbf{L})$

Let $* \in \{\text{SUM}, \text{MAX}, \infty\}$. The *optimum* associated with a weighted congestion game is defined by $\text{OPT}_* = \min_{\mathbf{P}} SC_*(\mathbf{P})$.

2.4 Nash Equilibria and Price of Anarchy

We are interested in a special class of (mixed) assignments called Nash equilibria [30, 31] that we describe here. Given a weighted congestion game and an associated mixed assignment \mathbf{P} , a player $i \in [n]$ is *satisfied* if it can not improve its private cost by unilaterally changing its strategy. Otherwise, player i is *unsatisfied*. The mixed assignment \mathbf{P} is a *Nash equilibrium* if and only if all players $i \in [n]$ are satisfied. Depending on the type of assignment, we differ between *pure*, *mixed* and *fully mixed* Nash equilibria.

The *mixed price of anarchy*, also called *coordination ratio* and denoted $\text{PoA}_{\text{mixed}}$, is the maximum value, over all instances Γ and Nash equilibria \mathbf{P} , of the ratio $\frac{SC_*(\mathbf{P})}{\text{OPT}_*}$. If we restrict to pure Nash equilibria, then we speak of the *pure price of anarchy* and denote it by PoA_{pure} .

2.5 Selfish Steps

Fix any pure assignment \mathbf{L} . In a *selfish step*, exactly one unsatisfied player is allowed to change its pure strategy such that its private cost decreases. A selfish step is *greedy* if the player chooses its best strategy. Clearly, selfish steps define a neighborhood of pure assignments that can be reached from \mathbf{L} . The assignment \mathbf{L} has an empty neighborhood if and only if \mathbf{L} is a Nash equilibrium. Thus, a pure Nash equilibrium corresponds to a local optimum. This stresses the close relationship of selfish steps on the one hand and local search processes on the other hand.

2.6 Special Weighted Congestion Games

Weighted Network Congestion Games. In a *weighted network congestion game* the strategies of a player correspond to paths from a source to a destination in a network. Thus, this class of games can be interpreted as *routing games*. If the players share the same source and destination, then we have a weighted *single-commodity* network congestion game, otherwise a weighted *multi-commodity* network congestion game. The underlying network of a weighted single-commodity network congestion game is

called *l-layered* if all paths from source to destination have length l .

KP-Model. Koutsoupias and Papadimitriou [21] considered a special weighted network congestion game, now widely known as the *KP-model*. In this model, each of the n players is allowed to use exactly one of m resources (here called *links*), that is, $S_i = [m]$ for all $i \in [n]$. The players are called *identical* if all weights are equal, otherwise *arbitrary*. Associated with each link $j \in [m]$ is a *capacity* c_j representing the rate at which link j processes *load*. Clearly, the latency on link j is $f_j(l_j) = \frac{l_j}{c_j}$, showing that the latency functions are linear. If c_1, \dots, c_m are equal, then the resources are *identical*, otherwise *related*. Denote $C = \sum_{j \in [m]} c_j$. In order to measure the social welfare of the system, Koutsoupias and Papadimitriou [21] considered the expected maximum latency.

A natural goal is to identify a Nash equilibrium with worst social cost for a given instance. For the model of related links, Gairing *et al.* [15] conjectured that, in case of its existence, the fully mixed Nash equilibrium is the worst Nash equilibrium with respect to social cost.

Fully Mixed Nash Equilibrium Conjecture ([15]). *Consider the model of arbitrary players and related links. Then, for any instance such that a fully mixed Nash equilibrium \mathbf{F} exists, and for any associated Nash equilibrium \mathbf{P} , $\text{SC}_\infty(\mathbf{P}) \leq \text{SC}_\infty(\mathbf{F})$.*

Routing Games on Parallel Links. We also consider variants of the KP-model to which we refer as *routing games on parallel links*. In particular, we investigate *restricted strategy sets* in which the players are only allowed to choose from a subset of links, that is, $S_i \subseteq [m]$ for all $i \in [n]$.

2.7 Exact Potential Games

A function $\Phi : (S_1 \times \dots \times S_n) \mapsto \mathbb{R}$ is an *exact potential function* for a game Γ if for every pure strategy profile $\mathbf{L} = (s_1, \dots, s_n)$, for every player $i \in [n]$ and for every strategy $s'_i \in S_i$, $\text{PC}_i(\mathbf{L}') - \text{PC}_i(\mathbf{L}) = \Phi(\mathbf{L}') - \Phi(\mathbf{L})$, where $\mathbf{L}' = (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$. In this case, Γ is an *exact potential game*. Since all exact potential games admit a pure Nash equilibrium (see e.g. [29]) these games are of interest in this paper.

3 Existence and Computation of Pure Nash Equilibria

Even though Nash was able to show that every finite game possesses a mixed Nash equilibrium, the question which class of games admits a pure Nash equilibrium remains open. In the case of its existence, it is of interest whether it is possible to compute a pure Nash equilibrium in polynomial time. In this section, we give some positive and some negative answers to both questions concerning the existence and the polynomial time computation. We start in Section 3.1 with routing games on parallel links and continue in Sections 3.2 and 3.3 with congestion games and weighted congestion games.

3.1 Routing Games on Parallel Links

We begin our survey with results on the KP-model. Afterwards we focus on games with restricted strategy sets.

KP-model. We first turn our attention to 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.* [10] showed that the LPT algorithm, first explored by Graham [16], yields some pure Nash equilibrium. Clearly, this holds for parallel links with *arbitrary non-decreasing* latency functions. For related links, the social cost of the Nash equilibrium computed by LPT approximates the social cost of an optimal assignment by a factor between 1.52 and 1.67 [12].

The second approach is to convert a given pure assignment into a Nash equilibrium without increasing the social cost. This conversion process is called *nashification*. Since selfish steps do not increase the social cost and any sequence of selfish steps eventually reaches a pure Nash equilibrium, selfish steps seem to be suitable for nashification. However, we have to use them carefully since the number of selfish steps may be exponential in the number of players before reaching a pure Nash equilibrium.

Theorem 1 ([7]). *Consider the model of arbitrary players and identical links. Then, there exists an instance and associated pure assignment for which the maximum length of a sequence of greedy selfish steps is at least*

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

Though there exist sequences of greedy selfish steps of exponential length, it is possible to use selfish steps to compute a Nash equilibrium in polynomial time if the links are identical. In particular, always moving an unsatisfied player with maximum weight to its best link requires at most n greedy selfish steps [15]. For related links, it is unknown whether selfish steps can be used to implement nashification in polynomial time. Feldmann *et al.* [9] chose a different approach not only based on selfish steps. Their algorithm relies on the following crucial observation.

Lemma 1 ([9]). *Consider the model of arbitrary players and related links. Then, for any pure assignment, a greedy selfish step of an unsatisfied player $i_1 \in [n]$ with weight w_{i_1} from a link $j_1 \in [m]$ to a link $j_2 \in [m]$ with $c_{j_1} \leq c_{j_2}$ makes no satisfied player $i_2 \in [n]$ with weight $w_{i_2} \geq w_{i_1}$ unsatisfied.*

The algorithm of Feldmann *et al.* [9] works in two phases. In the first phase, it fills up links with small capacities with players with small weight as close to $SC_{\text{MAX}}(\mathbf{L})$ as possible (but without exceeding $SC_{\text{MAX}}(\mathbf{L})$), and it collects all these users 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 the weights. Lemma 1 allows to show that this procedure results in a pure Nash equilibrium. Implementing the algorithm in a proper way, we get:

Theorem 2 ([9]). *Consider the model of arbitrary players and related links. Then, for any pure assignment \mathbf{L} , a pure Nash equilibrium \mathbf{L}' with $SC_{\infty}(\mathbf{L}') \leq SC_{\infty}(\mathbf{L})$ can be computed using $O(m^2n)$ time.*

Thus, we can apply the PTAS of Hochbaum and Shmoys [17] for scheduling *jobs* on related *machines* and then convert the computed assignment into a pure Nash equilibrium in polynomial time, and we get:

Corollary 1. *There is a PTAS for computing a best pure Nash equilibrium.*

Restricted Strategy Sets. Gairing *et al.* [13] considered a variant of the routing game on parallel links where there exists at least one player $i \in [n]$ with $S_i \subsetneq [m]$. So, the strategy sets of the players are *restricted*.

Gairing *et al.* [13] combined ideas from blocking flows and the generic PREFLOW-PUSH algorithm to derive a nashification algorithm for games with restricted strategy sets on identical links.

Theorem 3 ([13]). *Consider the model of arbitrary players with restricted strategy sets and identical links. Then, for any pure assignment \mathbf{L} , a pure Nash equilibrium \mathbf{L}' with $SC_\infty(\mathbf{L}') \leq SC_\infty(\mathbf{L})$ can be computed from \mathbf{L} using $O(rmA(\log W + m^2))$ time, where r is the number of distinct weights and $A = \sum_{i \in [n]} |S_i|$.*

Lenstra *et al.* [23] showed that an optimum assignment can be approximated within a factor of 2. It is worth mentioning that the nashification algorithm of Gairing *et al.* [13] improves this result since, for *any* given assignment \mathbf{L} , it computes a pure Nash equilibrium \mathbf{L}' with $SC_\infty(\mathbf{L}') \leq (2 - \frac{1}{w_1}) \cdot OPT_\infty$. Note that we can not hope to approximate an optimum assignment with factor less than $\frac{3}{2}$ unless $\mathcal{P} = \mathcal{NP}$ [23].

3.2 Congestion Games

In his seminal paper, Rosenthal [33] proved that $\Phi(\mathbf{L}) = \sum_{e \in E} \sum_{j=1}^{l_e(\mathbf{L})} f_e(j)$ is an exact potential function for congestion games. An immediate consequence follows:

Theorem 4 ([33]). *Every congestion game possesses a pure Nash equilibrium.*

Rosenthal's argumentation implies that every congestion game is an exact potential game. A result by Monderer and Shapley [29] shows that every exact potential game is closely related to a congestion game.

Theorem 5 ([29]). *Every finite exact potential game is isomorphic to a congestion game.*

Since every congestion game Γ possesses a pure Nash equilibrium the natural question arises whether it is possible to compute a pure Nash equilibrium for Γ in polynomial time. It is easy to see that this computational problem is in *PLS*. The class *PLS* (polynomial-time local search) introduced in [19] consists of local search problems for which local optimality can be verified in polynomial time. Many local search problems were shown to be complete for this class (see e.g. [19, 22, 37]), including graph partitioning, weighted satisfiability and traveling salesman problems. For none of these *PLS*-complete problems an algorithm is known that is able to compute a local optimum in polynomial time.

Using a sophisticated *PLS-reduction* Fabrikant *et al.* [8] proved that the computation of a pure Nash equilibrium for symmetric congestion games and asymmetric network congestion games is PLS-complete (see Figure 1). However, they showed that it is possible to calculate a pure Nash equilibrium for a symmetric network congestion game in polynomial time by using a min-cost flow algorithm.

	Symmetric	Asymmetric
Congestion Games	PLS-complete	PLS-complete
Network Congestion Games	Polynomial time	PLS-complete

Fig. 1. Complexity of computing pure Nash equilibria in congestion games [8]

We now switch to the class of *congestion games with player-specific payoff-functions* introduced by Milchtaich [28]. Here, a player always selects exactly one resource, that is, $S_1 = \dots = S_n = E$. Furthermore, the private cost of a player $i \in [n]$ on a resource $e \in E$ is described by a load dependent non-increasing latency function $f_e^i : \mathbb{R}^+ \mapsto \mathbb{R}^+$ that may be different from the latency function f_e^j for another player $j \neq i$. Milchtaich [28] considered these games with respect to pure Nash equilibria and sequences of selfish step. He showed:

Theorem 6 ([28]). *Every congestion game with player-specific payoff-function possesses a pure Nash equilibrium.*

Theorem 7 ([28]). *There exists a finite congestion game with player-specific payoff-function that admits a cycle of selfish steps, that is, a sequence of selfish steps starting and ending in the same assignment.*

It follows from the last theorem that games with player-specific payoff-functions do not admit an exact potential function.

3.3 Weighted Congestion Games

In this section we deal with weighted congestion games where the players may have different weights. Fotakis *et al.* [11] showed that there are such games that possess no pure Nash equilibrium. Moreover, they were able to prove that there is a subclass of games for which the existence of pure Nash equilibria is guaranteed.

Theorem 8 ([11]). *There exist instances of weighted single-commodity network congestion games for which there is no pure Nash equilibrium.*

Theorem 9 ([11]). *For any weighted multi-commodity network congestion game with linear latency functions, at least one pure Nash equilibrium exists.*

4 Price of Anarchy

The *mixed price of anarchy*, also known as *coordination ratio*, has been defined in the seminal work by Koutsoupias and Papadimitriou [21] as a measure of the extent to which non-cooperation approximates cooperation. Recall that it is defined as the worst-case ratio between the value of social cost in a Nash equilibrium and that of a social optimum. We present results on the pure and the mixed price of anarchy for routing games on parallel links in Section 4.1, for congestion games in Section 4.2, and for weighted congestion games in Section 4.3.

4.1 Routing Games on Parallel Links

We start with results on the KP-model. We then focus on the extension of this model to restricted strategy sets. Finally, we investigate routing games on parallel links with social cost defined as the sum of the private costs of the players.

KP-model. In the KP-model, latency functions are linear, social cost is defined as the expected maximum latency and the players may choose any link. For the case of identical links the pure price of anarchy is upper bounded by a constant. This does not hold for related links or mixed Nash equilibria. The bounds for mixed Nash equilibria are shown by first bounding the maximum expected load on a link and then applying a Hoeffding inequality [18]. All bounds are summarized in Figure 2.

	Pure Price of Anarchy	Mixed Price of Anarchy
Identical Links	$2 - \frac{2}{m+1}$ [15]	$\Theta\left(\frac{\log m}{\log \log m}\right)$ [5, 20]
Related Links	$\Theta\left(\frac{\log m}{\log \log m}\right)$ [5]	$\Theta\left(\frac{\log m}{\log \log \log m}\right)$ [5]

Fig. 2. Pure and mixed price of anarchy for the KP-model

Restricted Strategy Sets. In case of restricted strategy sets, even for identical links, the pure price of anarchy cannot be bounded by a constant. This also holds if the weights are identical. Figure 3 shows bounds on the pure price of anarchy. Note, that the bound for identical players and related links is only tight if $n = m$. Awerbuch *et al.* [2] further extended their result to mixed Nash equilibria.

Theorem 10 ([2]). *Consider the model of arbitrary players with restricted strategy sets and identical links. Then,*

$$\text{PoA}_{\text{mixed}} = \Theta\left(\frac{\log m}{\log \log \log m}\right).$$

Social Cost as Sum of Private Costs. Gairing *et al.* [14] considered another routing game on parallel links. In contrast to the KP-model, social cost is defined as the sum of the private costs of the players. This good natured definition of social cost makes the

	Identical Players	Arbitrary Players
Identical Links	$\Theta\left(\frac{\log m}{\log \log m}\right)$ [2, 13]	$\Theta\left(\frac{\log m}{\log \log m}\right)$ [2, 13]
Related Links	$O\left(\frac{\log n}{\log \log n}\right)$ [13]	$m - 1 \leq \text{PoA}_{\text{pure}} \leq m$ [13]

Fig. 3. Pure price of anarchy for the KP-model with restricted strategy sets

analysis significantly simpler and allows the investigation of general non-decreasing non-constant latency functions. For identical players, Gairing *et al.* [14] carried over an upper bound on the pure price of anarchy from the Wardrop-model [35] to the discrete setting.

Proposition 1 ([14]). *Consider the model of identical players and arbitrary links with non-decreasing and non-constant latency functions. If $xf_j(x) \leq \alpha \sum_{t=1}^x f_j(t)$ for all $x \in [n]$ and $j \in [m]$, then for any pure Nash equilibrium \mathbf{L} , $\text{SC}_{\text{SUM}}(\mathbf{L}) \leq \alpha \cdot \text{OPT}_{\text{SUM}}$.*

Corollary 2 ([14]). *Consider the model of identical players and arbitrary links. If the latency functions are polynomials with non-negative coefficients and maximum degree d , then the pure price of anarchy is bounded by $d + 1$.*

In case that all links have the same latency function $f(x) = x^d$, one can show the following bound on the mixed price of anarchy. B_k is the k 'th Bell number and counts the number of ways that a set of k elements can be partitioned into non-empty subsets.

Theorem 11 ([14]). *Consider the model of identical players and identical links with latency function $f(x) = x^d$, $d \in \mathbb{N}$. Then,*

$$\sup_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\text{SUM}}(\mathbf{P})}{\text{OPT}_{\text{SUM}}} = B_{d+1}.$$

4.2 Congestion Games

Recently, the pure price of anarchy found its way into congestion games [1, 4]. We restrict to results of Christodoulou and Koutsoupias [4] since only the abstract of the paper of Awerbuch *et al.* [1] was available (note that the latter paper also considers *weighted* congestion games). For congestion games with linear latency functions, Figure 4 summarizes results (both upper and lower bounds) on the pure price of anarchy. For the case of symmetric congestion games and social cost as the maximum of the private costs there is still a gap between the upper and the lower bound.

Christodoulou and Koutsoupias [4] also considered polynomial latency functions of degree d with non-negative coefficients. Figure 5 shows the corresponding bounds.

Both linear and polynomial latency functions were also considered in the Wardrop-model. Recall that in this model the social welfare of the system is defined as the sum of the edge latencies. The pure price of anarchy for linear latency functions is $\frac{4}{3}$ [35] whereas the pure price of anarchy for polynomial latency functions of degree d turned out to be $\Theta\left(\frac{d}{\log d}\right)$ [34].

	SC _{SUM}	SC _{MAX}
Symmetric	$\frac{5n-2}{2n+1}$	$\frac{5n-2}{2n+1} \leq \text{PoA}_{\text{pure}} \leq \frac{5}{2}$
Asymmetric	$\frac{5}{2}$	$\Theta(\sqrt{n})$

Fig. 4. Pure price of anarchy for congestion games with linear latency functions [4]

	SC _{SUM}	SC _{MAX}
Symmetric	$d^{\Theta(d)}$	$d^{\Theta(d)}$
Asymmetric	$d^{\Theta(d)}$	$\Omega(n^{d/(d+1)}), O(n)$

Fig. 5. Pure price of anarchy for congestion games with polynomial latency functions [4]

4.3 Weighted Congestion Games.

The mixed price of anarchy was also studied in weighted congestion games. Fotakis *et al.* [11] considered l -layered networks with identical edges each having the same linear latency function and social cost defined as the expected maximum latency.

Theorem 12 ([11]). *For weighted l -layered network congestion games with latency function $f_e(x) = x$ for all $e \in E$, the mixed price of anarchy for social cost as expected maximum latency is $O\left(\frac{\log m}{\log \log m}\right)$.*

This result is particularly interesting in comparison with the corresponding bound for the parallel link network (see Figure 2). It shows that under all l -layered networks the parallel link network has worst mixed price of anarchy.

5 Fully Mixed Nash Equilibria for Routing Games on Parallel Links

In routing games on parallel links, a *fully mixed* Nash equilibrium is a special Nash equilibrium, where each player chooses each link with strictly positive probability. Such a Nash equilibrium does not always exist. In this section, we give a characterization of instances with a fully mixed Nash equilibrium, we show its uniqueness and we study the Fully Mixed Nash Equilibrium Conjecture. We do this for two different routing games on parallel links. We would like to point out that there exist routing games on parallel links for which the Fully Mixed Nash Equilibrium Conjecture was disproved [24].

KP-model. Mavronicolas and Spirakis [26] were the first to consider fully mixed Nash equilibria. They showed for the KP-model, that if a fully mixed Nash equilibrium exists, it is unique and can be easily computed.

Theorem 13 ([26]). *Consider the model of arbitrary players and related links. Then, there exists a fully mixed Nash equilibrium \mathbf{F} if and only if*

$$f_{ij} = \left(1 - \frac{mc_j}{C}\right) \cdot \left(1 - \frac{W}{(n-1)w_i}\right) + \frac{c_j}{C} \in (0, 1)$$

for all $i \in [n]$ and $j \in [m]$. If \mathbf{F} exists, then \mathbf{F} is unique and $\mathbf{F} = (f_{ij})_{i \in [n], j \in [m]}$.

In particular, this implies that for the case of identical links the fully mixed Nash equilibrium uniquely exists and has probabilities $f_{ij} = \frac{1}{m}, \forall i \in [n], j \in [m]$.

In [15], the Fully Mixed Nash Equilibrium Conjecture was first explicitly stated for the KP-model, where social cost is defined as the expected maximum latency. Here, the ultimate settlement of this conjecture would reveal an interesting complexity-theoretical contrast between the worst-case pure and the worst-case mixed Nash equilibria. On the one hand, if the conjecture is valid, then the identification of the worst-case mixed Nash equilibrium is immediate in the cases where the fully mixed Nash equilibrium exists. On the other hand, Gairing *et al.* [15] showed that the worst-case pure Nash equilibrium is not $(2 - \frac{2}{m+1} - \varepsilon)$ -approximable even on identical links.

Theorem 14 ([15]). *Consider the model of arbitrary players and identical links. If, for any ε with $0 < \varepsilon \leq 1 - \frac{2}{m+1}$, the worst-case pure Nash equilibrium is $(2 - \frac{2}{m+1} - \varepsilon)$ -approximable, then $\mathcal{P} = \mathcal{NP}$.*

This result also unfolds an interesting contrast between best and worst-case pure Nash equilibria. For any $\varepsilon > 0$, a pure Nash equilibrium \mathbf{L} with $\text{SC}_\infty(\mathbf{L}) \leq (1 + \varepsilon) \cdot \text{OPT}_\infty$ can be computed in polynomial time whereas the computation of a pure Nash equilibrium \mathbf{L}' with $\text{SC}_\infty(\mathbf{L}') \geq (1 + \varepsilon) \cdot \text{OPT}_\infty$ is \mathcal{NP} -hard.

So far, the Fully Mixed Nash Equilibrium Conjecture has been proved only for some special cases, namely, two players on identical links [15], two identical players on related links and identical players on two identical links [24]. Furthermore, it was shown up to a factor of 49.02 in case of identical players and related links [10] and up to a factor of $2h(1 + \varepsilon)$ for arbitrary players on identical links, if $n = m$ sufficient large [15], where h is the factor between the maximum and the average weight of the players.

On the other hand, Gairing *et al.* [15] proved that the private costs of all players in a Nash equilibrium are upper bounded by their private costs in the fully mixed Nash equilibrium. This directly implies:

Theorem 15 ([15]). *Consider the model of arbitrary players and related links. If the fully mixed Nash equilibrium \mathbf{F} exists, then, for any mixed Nash equilibrium \mathbf{P} , we have $\text{SC}_{\text{SUM}}(\mathbf{P}) \leq \text{SC}_{\text{SUM}}(\mathbf{F})$ and $\text{SC}_{\text{MAX}}(\mathbf{P}) \leq \text{SC}_{\text{MAX}}(\mathbf{F})$.*

Social Cost as Sum of Private Costs. Fully mixed Nash equilibria were also considered for identical players and general non-decreasing and non-constant latency functions with respect to social cost defined as the sum of the private costs [14]. In order to characterize instances where the fully mixed Nash equilibrium exists, Gairing *et al.* [14] introduced two classes of links, namely *dead links* and *special links*. They showed that in any Nash equilibrium, none of the players is assigned to a dead link. Moreover, there exists at most one player who is assigned to any of the special links. Availing these results, they could give the following thorough characterization.

Theorem 16 ([14]). *Consider the model of identical players and links with non-decreasing and non-constant latency functions. Then, there exists a fully mixed Nash equilibrium \mathbf{F} if and only if there are no special and no dead links. If \mathbf{F} exists then \mathbf{F} is unique.*

For every instance define the *generalized fully mixed Nash equilibrium* as the fully mixed Nash equilibrium for the instance where the links are restricted to non-special and non-dead links. If latency functions are non-decreasing, non-constant and *convex*, then one can show, that the private cost of each player in a Nash equilibrium is upper bounded by its private cost in the generalized fully mixed Nash equilibrium. This directly implies:

Theorem 17 ([14]). *Consider the model of identical players and links with non-decreasing, non-constant and convex latency functions. Then, for any Nash equilibrium \mathbf{P} and generalized fully mixed Nash equilibrium \mathbf{F} , $SC_{\text{SUM}}(\mathbf{P}) \leq SC_{\text{SUM}}(\mathbf{F})$ and $SC_{\text{MAX}}(\mathbf{P}) \leq SC_{\text{MAX}}(\mathbf{F})$.*

6 Open Problems

The flourishing interest in weighted congestion games resulted in a multitude of results and methods, but raised even more questions remaining tantalizingly open. We only state some of them:

- Although the results of Nash [30, 31] guarantee the existence of a Nash equilibrium in strategic games, the computational complexity of computing a Nash equilibrium is open even if only two players are involved.
- Which classes of symmetric weighted network congestions games possess a pure Nash equilibrium? For which classes is it possible to compute such a pure Nash equilibrium in polynomial time?
- It is impossible to approximate a worst-case pure Nash equilibrium within a factor better than $2 - \frac{2}{m+1}$ in the KP-model with identical links [15]. To which extent is it possible to approximate a worst-case pure Nash equilibrium in the KP-model with related links or in more general settings?
- Most of the known bounds on the price of anarchy for network congestion games were shown with respect to social cost defined as sum or maximum of the private costs of the players [1, 4]. What is the price of anarchy if social cost is defined as expected maximum latency?
- For the KP-model, Gairing *et al.* [15] showed that the private costs of all players in a Nash equilibrium are bounded from above by their private costs in the fully mixed Nash equilibrium. For which classes of network congestion games does this property still hold?
- If the players are identical and the links are related, then the Fully Mixed Nash Equilibrium Conjecture holds up to a factor of $2h(1 + \varepsilon)$, where h is the factor between the maximum and the average weight of the players [15]. Does there exist an approximation factor independent of h ?

References

1. 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)*, 2005.
2. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. In *Proceedings of the 1st International Workshop on Approximation and Online Algorithms (WAOA'03)*, LNCS 2909, pages 41–52, 2003.
3. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
4. 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)*, 2005.
5. 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.
6. S. C. Dafermos and F. T. Sparrow. The Traffic Assignment Problem for a General Network. *Journal of Research of the National Bureau of Standards, Series B*, 73(2):91–118, 1969.
7. E. Even-Dar, A. Kesselmann, and Y. Mansour. Convergence Time to Nash Equilibria. In *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, LNCS 2719, pages 502–513, 2003.
8. 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.
9. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the Coordination Ratio for a Selfish Routing Game. In *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, LNCS 2719, pages 514–526, 2003.
10. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP'02)*, LNCS 2380, pages 123–134, 2002.
11. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. Accepted to *Theoretical Computer Science*.
12. D. K. Friesen. Tighter Bounds for LPT Scheduling on Uniform Processors. *SIAM Journal on Computing*, 16(3):554–560, 1987.
13. 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.
14. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, LNCS 3142, pages 645–657, 2004.
15. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme Nash Equilibria. In *Proceedings of the 8th Italian Conference on Theoretical Computer Science (ICTCS'03)*, LNCS 2841, pages 1–20, 2003. Also accepted to *Theoretical Computer Science*, Special Issue on *Game Theory Meets Theoretical Computer Science*.
16. R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.
17. 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.

18. W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *American Statistical Association Journal*, 58(301):12–30, 1963.
19. D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How Easy is Local Search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.
20. E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
21. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, LNCS 1563, pages 404–413, 1999.
22. M. W. Krentel. On Finding and Verifying Locally Optimal Solutions. *SIAM Journal of Computing*, 19(4):742–729, 1990.
23. J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation Algorithms for Scheduling Unrelated Parallel Machines. *Mathematical Programming*, 46:259–271, 1990.
24. T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash Equilibrium? In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, LNCS 2747, pages 551–561, 2003.
25. A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
26. 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.
27. R. D. McKelvey and A. McLennan. Computation of Equilibria in Finite Games. In *Handbook of Computational Economics*, 1996.
28. I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
29. D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
30. 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.
31. J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
32. 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.
33. R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
34. T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
35. T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
36. T. Roughgarden and É. Tardos. Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games. *Games and Economic Behaviour*, 47(2):389–403, 2004.
37. A. A. Schäffer and M. Yannakakis. Simple Local Search Problems that are Hard to Solve. *SIAM Journal of Computing*, 20(1):56–87, 1991.
38. V. Vazirani. *Approximation Algorithms*. Springer Verlag, 2001.
39. J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.
40. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
41. 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, 1956.