Price of anarchy for parallel link networks with generalized mean objective

Pieter Kleer[1]

Abstract

The price of anarchy is the most well-known measure for quantifying the inefficiency of equilibrium flows in traffic networks and routing games. In this work, we give unifying price of anarchy bounds for atomic and non-atomic parallel link routing games with polynomial cost functions under various cost objectives including the arithmetic mean, geometric mean and worst-case cost objective. We do this through the study of the generalized p-mean as cost objective, and obtain upper bounds on the price of anarchy in terms of this parameter p. Our bounds unify existing results from the literature, and, in particular, give alternative proofs for price of anarchy results in parallel link routing games with polynomial cost functions under the geometric mean objective obtained by Vinci et al. (ACM Trans Econ Comput 10(2):41, 2022). We recover those simply as a limiting case. To the best of our knowledge, these are the first price of anarchy bounds that capture multiple cost objectives simultaneously in a closed-form expression.

Keywords Nash equilibrium · Wardrop flow · Price of anarchy · Generalized mean · Parallel link networks

1 Introduction

One of the fundamental problems at the intersection of transportation research and game theory is to quantify the inefficiency of equilibrium flows in traffic networks. Equilibrium flows, also known as Wardrop flows or Nash equilibria, arise as a result of autonomous users acting strategically to minimize their individual travel time. These flows often do not correpond to a system optimum, which is desired from a network designer's point of view. Pigou (1920) was the first to (theoretically) demonstrate the possibility of inefficiency, and ever since, this topic has been of great

✉ Pieter Kleer
  p.s.kleer@tilburguniversity.edu

1 Department of Econometrics and Operations Research, Tilburg University, Warandelaan 2,
   5037 AB Tilburg, The Netherlands

interest to both researchers and practitioners in various areas, see, e.g., Dubey (1986), Roughgarden (2005), Braess et al. (2005).

Nowadays, the most well-known notion to quantify the inefficiency of equilibrium flows is the *price of anarchy*, introduced by Koutsoupias and Papadimitriou (1999, 2009). The price of anarchy has been studied extensively in the last 20 years in various routing models, both from a theoretical perspective, see, e.g., Roughgarden and Tardos (2002), Roughgarden (2003), Christodoulou and Koutsoupias (2005), Aland et al. (2011), Koutsoupias and Papadimitriou (1999), Caragiannis et al. (2011), Roughgarden (2015), as well as in practice, see, e.g., Monnot et al. (2017), Zhang et al. (2018), Benita et al. (2020). It is defined as the ratio between the *quality* of a (worst-case) equilibrium flow and a system optimum. The quality is measured with respect to a cost objective, such as the artithmetic mean (i.e. average) user cost, called the *utilitarian* objective, or the worst-case cost of any individual user, called the *egalitarian* objective. Recently, also the geometric mean of user costs has gained attention (Vinci et al. 2022). In the context of welfare maximization problems, the geometric mean of user costs is known as *Nash social welfare* (John 1950) in game theory and economics, where it has been studied extensively. However, much less is known about this objective in the context of cost minimization problems, which also was one of the motivations for studying this objective in Vinci et al. (2022). One appealing property of the geometric mean is that the price of anarchy is invariant with respect to rescaling the costs of individual users. This means users can measure their costs with respect to different units such as time or money.[1]

The *choice* of cost objective, which is the topic of this work, plays an important role in how large the price of anarchy of equilibrium flows can be. For example, an equilibrium flow that might be considered "inefficient" under the utilitarian cost objective, can be efficient under the egalitarian objective as we will see in Example 1.1. The choice of cost objective is therefore also important for practitioners, who have to decide whether it is necessary to take measures to reduce the inefficiency, such as introducing road tolls.

Example 1.1 below can be seen as a discrete version of the well-known example of Pigou (1920) that first illustrated inefficieny in traffic networks. The atomic routing model used in Example 1.1 is formalized in Sect. 2. We study the three cost objectives mentioned above: the utilitarian, egalitarian and geometric mean objective.

Example 1.1 We consider an (atomic) routing game played on a parallel link network consisting of two links T and B, as illustrated in Fig. 1, that both are equipped with a cost (or latency) function $c_r(x)$ for $r \in \{T, B\}$. The game has two users, a and b, both represented by a unit of flow. Both have to choose one of the two links in the network over which they have to route their unsplittable unit of flow.

---

[1] The scaling invariance follows from the formulation in (1.2).

Fig. 1 Parallel link network consisting of two arcs

$$c_T(x) = \tfrac{1}{2}x$$

$$c_B(x) = 1$$
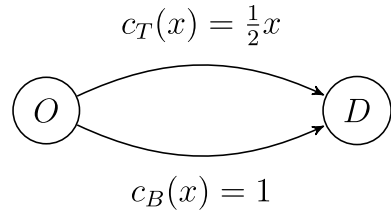
An assignment of users to links is called a strategy profile. For example, if both users choose link T, we have the strategy profile $s = (s_a, s_b) = (T, T)$. The cost $C_i(s)$ of player i in a strategy profile s is the cost incurred on the chosen links, which is $c_r(x_r)$ where x_r is the the total amount of flow on link r. In the strategy profile $s = (T, T)$ this means both a and b have cost $C_a(s) = C_b(s) = c_T(2) = \frac{1}{2} \cdot 2 = 1$. A strategy profile is called a (pure) Nash equilibrium if no user can improve its cost by switching to the other link. The profile $s = (T, T)$ is then indeed a Nash equilibrium: If one of the users would deviate to link B, their new cost will equal one, which is the same as they have in the current profile s.

We next describe three possible cost objectives $C(\cdot)$ that can be used to quantify the inefficiency of the Nash equilibrium s. For all these objectives the profile $s^* = (s_a^*, s_b^*) = (T, B)$ minimizes the cost objective over all strategy profiles.

- Arithmetic mean cost (utilitarian objective): It holds that $C(s^*) = (C_a(s^*) + C_b(s^*))/2 = 0.75$. The Nash equilibrium $s = (2, 2)$ has $C(s) = (1 + 1)/2 = 1$. This means the inefficiency of this Nash equilibrium is $C(s)/C(s^*) = 4/3$.
- Worst-case cost (egalitarian objective): It holds that $C(s^*) = \max\{C_a(s^*), C_b(s^*)\} = \max\{1, 1/2\} = 1$. The Nash equilibrium has $C(s) = \max\{1, 1\} = 1$. This means that the inefficiency is $C(s)/C(s^*) = 1$.
- Geometric mean cost: It holds that $C(s^*) = \left(C_a(s^*)C_b(s^*)\right)^{\frac{1}{2}} = (1 \cdot 1/2)^{\frac{1}{2}} = 1/\sqrt{2}$. The Nash equilibrium has $C(s) = (1 \cdot 1)^{\frac{1}{2}} = 1$. This means the inefficiency is $C(s)/C(s^*) = \sqrt{2}$.

As the above example illustrates, the choice of cost function indeed plays an important rule in how large the inefficiency of an equilibrium flow can be. In the case of the egalitarian objective, there is no inefficiency, whereas for the utilitarian and geometric mean objective, the inefficiency is roughly 33% and 41%, respectively. These observations give rise to the following, somewhat informal, question:

Is it possible to quantify the price of anarchy for different cost objectives in one formula?

The main contribution of this work is that we answer the above question affirmatively for *parallel link* routing games with cost functions on the links being

polynomials of degree at most d by considering the so-called *generalized p-mean cost objective* (defined in Sect. 1.1). We study both the atomic routing model of Rosenthal (1973) as well as the non-atomic routing model of Wardrop (1952). To the best of our knowledge, the results presented here are the first inefficiency results for routing games that capture multiple cost objectives such as the utilitarian, egalitarian and geometric mean objective at once. We remark that, although parallel link networks are the simplest type of traffic networks, they have been studied extensively in the areas of algorithmic game theory and operations research, see, e.g., Caragiannis (2008), Gairing et al. (2008), Vinci et al. (2022), Colini-Baldeschi et al. (2019), Bonifaci et al. (2011), Harks et al. (2019). Adding to their significance, they can also naturally be interpreted as scheduling, or load balancing, problems where the links play the roles of machines and users control jobs that have to be scheduled. See, e.g., Caragiannis (2008) and references therein.

In the utilitarian case, it is well-known that the price of anarchy typically can be characterized by the type of cost functions the links of the traffic network are allowed to have (e.g., polynomials of degree at most $d = 1, 2, 3, \ldots$). This was first established in the seminal work of Roughgarden (2003); see also his *smoothness* framework (Roughgarden 2015) and Sect. 1.3 for more related work. Polynomial cost functions also have practical relevance, e.g., the cost functions proposed by the Bureau of Public Roads (States 1964) for modelling congestion in road networks are polynomials of degree $d = 4$.

1.1 Our contributions

Parallel link routing games[2] consist of a set $R = \{1, \ldots, m\}$ of parallel links between an origin O and destination D, and a set of users $N = \{1, \ldots, n\}$. Every user $i \in N$ has a subset of links $S_i \subseteq R$, called its *strategy set*, that they are allowed to use. Every link is equipped with a nonnegative, nondecreasing cost function $c_r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. The cost $C_i(s)$ of a player in strategy profile $s \in \times_i S_i$ is the cost incurred on the chosen link. To get a general framework capturing all the different cost objectives, we consider the so-called *generalized p-mean objective*

$$C_w^p(s) = \left(\frac{1}{\sum_{i \in N} w_i} \sum_{i \in N} w_i \cdot C_i(s)^p \right)^{\frac{1}{p}} \tag{1.1}$$

for $p \in (0, \infty)$, where $w = (w_1, \ldots, w_n)$ is indicating for every $i \in N$ the amount of (unsplittable) flow w_i that user i has to route over one of its links. If $w_i = 1$ for all $i \in N$, the game is called *unweighted*. Furthermore, we have

$$C_w^0(s) := \lim_{p \to 0} C_w^p(s) = \left(\prod_{i \in N} C_i(s)^{w_i} \right)^{1/\sum_{i \in N} w_i} \quad \text{and} \quad C_w^\infty(s) := \lim_{p \to \infty} C_w^p(s) = \max_{i \in N} C_i(s). \tag{1.2}$$

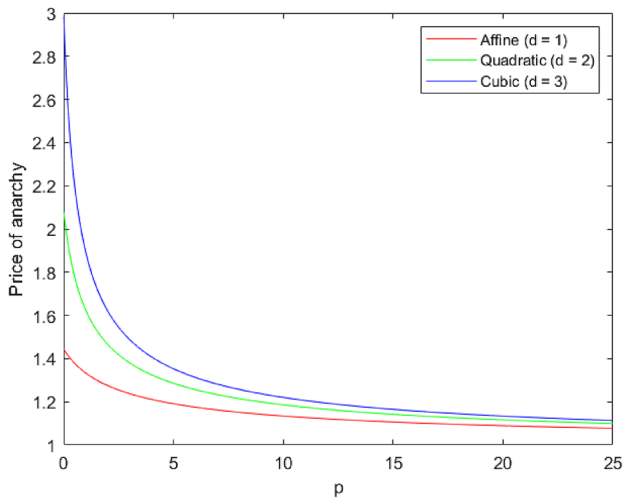[2] Formal definitions are given in Sect. 2.

Fig. 2 Price of anarchy bounds under the generalized p-mean objective for non-atomic parallel link routing games for $d = 1, 2, 3$ (bottom, middle and top function, respectively)

The generalized p-mean indeed captures the utilitarian objective ($p = 1$), egalitarian objective ($p = \infty$) and geometric mean objectives ($p = 0$) as special cases, as well as p-norm objectives for $p \geq 1$.

We next summarize our results for parallel link routing games with cost functions being polynomials of degree at most $d \in \mathbb{N}$. We emphasize that we obtain a price of anarchy bound for all values of $p \in [0, \infty]$, which allows us to also go beyond the utilitarian, egalitarian and geometric mean cost objectives. A comparison with existing work, that we capture as special cases, is given in Table 1. We consider both the atomic model of Rosenthal (1973) and the non-atomic model of Wardrop (1952), whose difference, roughly speaking, lies in whether it is assumed that the w_i are positive constants, or whether there is a continuum of players all controlling an infinitesimally small amount of flow (i.e., $w_i \to 0$). For all models, we first obtain a bound for $p \in (0, \infty)$, after which we obtain price of anarchy bounds for $p \in \{0, \infty\}$ simply as limiting cases, using the expressions in (1.2). In particular, for $p = 0$, this yields an alternative derivation of the price of anarchy results of Vinci et al. (2022) for games with polynomial cost functions and geometric mean objective, which is our main technical contribution.

- *Non-atomic model:* This is the model with a continuum of users that all control an infinitesimally small fraction of flow.[3] We show in Theorem 3.2 that, for $p \in (0, \infty)$, the price of anarchy is upper bounded by

---

[3] One may think of this model as $w_i \to 0$ for all $i \in N$ with $|N| \to \infty$, in such a way that the total amount of flow in the network remains constant.

$$\left(1 - \frac{dp}{(dp + 1)^{(dp+1)/(dp)}}\right)^{-\frac{1}{p}}. \tag{1.3}$$

As limiting cases we get that for $p = 0$ the price of anarchy is upper bounded by $(e^{1/e})^d$, and for $p = \infty$ by 1. These bounds also hold for atomic symmetric unweighted parallel link routing games, as explained in Remark 3.5. For an illustration of the obtained bounds in this case, see Fig. 2.

- *Weighted atomic model:* This is the general model as described above where user $i \in N$ has to route w_i units of (unsplittable) flow over one of its links. We show in Theorem 3.6 that, for $p \in (0, \infty)$, the price of anarchy is upper bounded by

$$(\Phi_{pd} + 1)^d, \tag{1.4}$$

where Φ_{pd} is the unique nonnegative solution to $(x + 1)^{pd} = x^{pd+1}$. As limiting cases, we obtain that for $p = 0$ the price of anarchy can be upper bounded by 2^d, but it might be unbounded for $p \to \infty$.

- *Unweighted atomic model:* This is the above model with $w_i = 1$ for all $i \in N$. We show in Theorem 3.9 that, for $p \in (0, \infty)$, the price of anarchy is upper bounded by

$$\left(\frac{(k + 1)^{2pd+1} - k^{pd+1}(k + 2)^{pd}}{(k + 1)^{pd+1} - (k + 2)^{pd} + (k + 1)^{pd} - k^{pd+1}}\right)^{\frac{1}{p}}, \tag{1.5}$$

where $k = \lfloor \Phi_{pd} \rfloor$ and Φ_{pd} is again the unique nonnegative solution to $(x + 1)^{pd} = x^{pd+1}$. Again, as limiting cases, we obtain that for $p = 0$ the price of anarchy can be upper bounded by 2^d, but it might be unbounded for $p \to \infty$.

The bounds we obtain in (1.3)–(1.5) resemble those for the case $p = 1$ in Table 1. This follows from the fact that we derive our bounds by appropriately adjusting the proof techniques for the respective setting with $p = 1$. We do this by showing, roughly speaking, that it suffices to study the utilitarian objective in the game where the cost functions $c_r(x)$ are replaced by $c_r(x)^p$, i.e., the function c_r raised to the power p. For the non-atomic case, this can be done by using the results of Correa et al. (2004), but for the atomic cases this requires non-trivial adjustments of the proofs by Aland et al. (2011). One of the technical tools we use for this is *Descartes' rule of signs* (see Appendix 3), which allows one to say something about the number of nonnegative roots of a polynomial.

It is interesting to note that the bounds obtained by Vinci et al. (2022) for the geometric mean objective have a much easier form than their counterparts for the arithmetic mean cost objective. This also motivates the study of the price of anarchy of the geometric mean objective in other classes of (routing) games.

1.2 Discussion and lower bounds

It is interesting to note that there is a clear contrast in how the obtained price of anarchy bounds behave when comparing the non-atomic and atomic routing models.

Table 1 Overview of special cases of the upper bounds in (1.3)–(1.5) for the price of anarchy in parallel link routing games for the geometric mean, utilitarian and egalitarian objectives

Model	$p = 0$	$p = 1$	$p = \infty$
Non-atomic	$\left(e^{1/e}\right)^d$ (Vinci et al. 2022)	$\left(1 - \frac{d}{(d+1)^{(d+1)/d}}\right)^{-1} \approx \Theta(d/\log(d))$ (Correa et al. 2004)	1
Weighted atomic	2^d (Vinci et al. 2022)	$(\Phi_d + 1)^d \approx \Theta(d/\log(d))^{d+1}$ (Aland et al. 2011)	∞ (Gairing et al. 2008)
Unweighted atomic	2^d (Vinci et al. 2022)	$\frac{(k+1)^{2d+1} - k^{d+1} - (k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}} \approx \Theta(d/\log(d))^{d+1}$ (Aland et al. 2011)	∞ (Gairing et al. 2008)

Here Φ_d is the unique nonnegative solution to $(x + 1)^d = x^{d+1}$ and $k = \lfloor \Phi_d \rfloor$. (The notation $f(x) = \Theta(g(x))$ means that there exist constants k_1 and k_2 such that $k_1 g(x) \le f(x) \le k_2 g(x)$ for x large enough)

Whereas for the non-atomic model, the price of anarchy is a decreasing function of p (see Fig. 2), it is increasing for both the unweighted and weighted atomic model. This also shows that not only the choice of cost objective plays a large role, but also the underlying model being considered.

Studying the generalized p-mean objective on more general networks is a very interesting direction for future research. One of the technical issues that arises is that the reduction sketched at the end of Sect. 1.1 breaks down beyond parallel link networks, and various nonlinearity issues arise. In fact, Vinci et al. (2022) show that there exists a congestion game[4] with affine cost functions where the inefficiency of a pure Nash equilibrium under the geometric mean objective can grow with the number of players of the game, which stands in stark contrast to the utilitarian setting, for which the price of anarchy is known to be at most 5/2 (Christodoulou and Koutsoupias 2005) in congestion games with affine cost functions.

Lower bounds. We note that the bounds we obtain in (1.3)–(1.5), and the special cases in Table 1, are known to be tight for various values of p and d. The non-atomic bound for $p \in (0, \infty)$ in Theorem 3.2 can be shown to be tight by using a similar construction as Correa et al. (2004) for polynomials of degree d (see Appendix 4). For the case $p = 0$, tightness is shown in Vinci et al. (2022).

For (symmetric) weighted atomic parallel link games, tightness for $p = 1$ is shown in Bhawalkar et al. (2014) for any d. A more general lower bound, using only mild assumptions on the cost functions, is provided in Bilò and Vinci (2017). Similar constructions and reasonings as in Bhawalkar et al. (2014), Bilò and Vinci (2017) seem to be able to give a tight lower bound for an arbitrary p. Indeed, Bilò and Vinci (2017, Theorem 1) should yield tightness of the bound in (1.4) for symmetric weighted atomic parallel link games; we leave the details of this to the interested reader.[5] Tightness for $p = 0$ is shown in Vinci et al. (2022).

For unweighted atomic parallel link games, with $p = \infty$, Gairing et al. (2008) show that the price of anarchy grows with the number of players (or links) already in games with affine cost functions ($d = 1$), which leads to a best possible upper bound of ∞ in case the price of anarchy is parameterized by the cost functions of the links. This explains the "∞" price of anarchy bounds for $p = \infty$ in Table 1. Tightness of the bound in (1.5) for the case $d = 1$ and $p = 1$ is shown in Caragiannis et al. (2011), and for every d in the case of general network routing games, where the topology does not have to be a set of parallel links, still with $p = 1$, in Aland et al. (2011). The lower bound in Caragiannis et al. (2011) is generalized in Gairing and Schoppmann (2007) to arbitrary polynomical cost functions of degree at most d. Similary as for weighted games, using (Bilò and Vinci 2017, Theorem 3) this time, it should be possible to show that the bound in (1.5) is tight.

---

[4] Congestion games form a generalization of network routing games.

[5] The result in Bilò and Vinci (2017, Theorem 1) works in a blackbox fashion, implicity showing that if smoothness parameters (λ, μ), as in Roughgarden's smoothness framework (Roughgarden 2015), are best possible, then one automatically gets tightness of the obtained price of anarchy bound. This argument may also be applied to the results obtained in Sect. 3 assuming the parameters (λ, μ) obtained there are optimal. This comment also applies to unweighted atomic parallel link games; see Bilò and Vinci (2017, Theorem 3). We leave the details to the interested reader.

1.3 Further related work

The price of anarchy in general (network) routing and congestion games has been studied extensively in both the non-atomic and atomic setting. In one of the seminal works on the price of anarchy for routing games, Roughgarden and Tardos (2002) showed that the price of anarchy for affine cost functions equals 4/3 in the non-atomic case. Tight bounds for general classes of cost functions, including polynomial cost functions, were obtained by Roughgarden (2003) and Correa et al. (2004). In the atomic unweighted setting, Christodoulou and Koutsoupias (2005) first showed a bound of 5/2 for games with affine cost functions. Tight bounds for polynomial cost functions, both in the unweighted and weighted setting, were obtained by Aland et al. (2011). These results, applied to the special case of parallel link networks, are captured by our formulas in (1.3)–(1.5). Further works considering congestion games with polynomial cost functions are, e.g., Gairing et al. (2006), Gairing and Schoppmann (2007), Christodoulou et al. (2011), Bhawalkar et al. (2014), Bilò and Vinci (2017), Bilò and Vinci (2019).

Various extensions and variations on the classical models of Rosenthal (1973) and Wardrop (1952) have been proposed in the literature, which, from a technical perspective, can often be seen as quantifying the price of anarchy along one or more parameters induced by the extended setting of interest (similar to what we do in this work by means of the parameter p). Such models are used to capture, e.g., altruistic behavior (Chen et al. 2014; Schröder 2020), risk-aversion (Fotakis et al. 2015), tolls (or taxes) (Caragiannis et al. 2010; Bonifaci et al. 2011) or uncertainty (Piliouras et al. 2016; Cominetti et al. 2019). A unifying framework of such models in the case of affine cost functions can be found in Kleer and Schäfer (2019b), where the (utilitarian) cost objective, as well as the individual user costs, contain a parameter along which the price of anarchy is studied. There is also a line of work quantifying the impact on the quality of equilibrium flows, as a result of perturbation or uncertainties in the input of the game, where instead of comparing to an optimal flow, one instead compares the new (perturbed) equilibrium flow to the original (unperturbed) equilibrium flow. This leads to new inefficiency notions such as (in chronological order of introduction) the *price of risk aversion* (Nikolova and Stier-Moses 2014; Lianeas et al. 2019), the *deviation ratio* (Kleer and Schäfer 2017, 2019a) or the *price of satisficing* (Takalloo and Kwon 2020). Inefficiency results for such notions are typically also parameterized by an input parameter.

Another fundamental inefficiency notion that has been studied extensively is the *price of stability* (Anshelevich et al. 2008) that compares the best possible social cost of a Nash equilibrium with that of a social optimum, i.e., the minimal "price" a system designer has to pay in order to get a system which is stable from the players' perspective. Further references here are, e.g., Christodoulou and Gairing (2015), Christodoulou et al. (2019), Kleer and Schäfer (2021).

1.4 Outline

We continue in Sect. 2 with a formal description of the routing models we consider in this work. After that, we prove our upper bounds on the price of anarchy under the generalized p-mean objective in Sect. 3. We conclude in Sect. 4.

2 Routing models

In this section we describe the routing models studied in this work together with the price of anarchy under the generalized p-mean objective.

2.1 Atomic parallel link routing game

A *weighted atomic parallel link routing game* $\Gamma = (N, R, (w_i)_{i \in N}, (C_i)_{i \in N})$ is given by a set of users $N = \{1, \ldots, n\}$ and a set of parallel (directed) links $R = \{r_1, \ldots, r_m\}$ between an origin O and destination D (as in Fig. 1). Every user i has a strategy set $\mathcal{S}_i \subseteq R$ of links it is allowed to use, and a weight $w_i \geq 0$. In case $\mathcal{S}_i = R$ for every $i \in N$, the game is called *symmetric*.[6] Every link $r \in R$ is equipped with a cost function $c_r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that is assumed to be non-decreasing (and non-negative by definition). User i places its (unsplittable) flow w_i on one of its links $r \in \mathcal{S}_i$. That is, if in strategy profile s user i chooses link s_i, we write

$$x_j^i(s) = \begin{cases} w_i & \text{if } j = s_i \\ 0 & \text{otherwise.} \end{cases}$$

The *(total) flow* on link $r \in R$ is defined as $x_r(s) = \sum_{i \in N} x_r^i(s)$. For convenience, we sometimes write x_r instead of $x_r(s)$ if it's clear what strategy profile we are considering. The cost of user i on its chosen link s_i is given by

$$C_i(s) = c_{s_i}(x_{s_i}).$$

[7]

Note that all users placing their weight on the same link, have the same cost.

We write $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ for the strategy profile in which s_i is left out. Slightly abusing notation, we then write $(s_i', s_{-i}) = (s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n)$. A strategy profile $s \in \times_i \mathcal{S}_i$ is called *pure Nash equilibrium* if

$$C_i(s) \leq C_i(s_i', s_{-i}) \tag{2.1}$$

---

[6] What we call a symmetric parallel link routing game is sometimes referred to as a parallel link routing game in the literature. What we call a (general) parallel link routing game is sometimes referred to as a restricted parallel link routing game in the literature.

[7] We note that sometimes the cost of a user is defined as $w_i c_r(x_r(s))$ but this is equivalent to our definition as the weight is taken into account in the objective C_w^p (and leaving it out does not affect the set of pure Nash equilibria).

for all $i \in N$ and unilateral deviations $s_i' \in S_i$. We write PNE(Γ) for the set of all pure Nash equilibria of the game Γ.

We define the *p-price of anarchy (p-PoA)* as

$$p\text{-PoA}(\Gamma) = \frac{\max_{s \in \text{PNE}(\Gamma)} C_w^p(s)}{\min_{t \in \times_i S_i} C_w^p(t)} \tag{2.2}$$

with $C_w^p(s)$ as in (1.1) for $p \in (0, \infty)$ or (1.2) for $p \in \{0, \infty\}$. It is important to note that the set PNE(Γ) is the same for every value of $p \in [0, \infty]$, as the inequality in (2.1) does not depend on p. A strategy minimizing the cost objective C_w^p is called a *socially optimal strategy profile (or system optimum)*. We next argue that in order to study the *p*-PoA, it suffices to study the utilitarian case in the game where the cost functions are raised to the power p.

For $p \in (0, \infty)$, in order to derive an upper bound of $\alpha \geq 1$ on the *p*-PoA, it suffices to show

$$C_w^p(s) \leq \alpha \cdot C_w^p(s^*)$$

where $s = (s_1, \ldots, s_n)$ is any pure Nash equilibrium and $s^* = (s_1^*, \ldots, s_n^*)$ a socially optimal strategy profile. This is equivalent to showing

$$\sum_{i \in N} w_i \cdot C_i(s)^p \leq \alpha^p \sum_{i \in N} w_i \cdot C_i(s^*)^p. \tag{2.3}$$

As every user chooses one link, this reduces to

$$\sum_{i \in N} w_i \cdot c_{s_i}(x_{s_i})^p \leq \alpha^p \sum_{i \in N} w_i \cdot c_{s_i^*}(x_{s_i^*})^p \tag{2.4}$$

Note that for any strategy profile $t = (t_1, \ldots, t_n)$, we may write the cost objective as an aggregation over link contributions,[8] i.e.,

$$\sum_{i \in N} w_i \cdot c_{t_i}(x_{t_i})^p = \sum_{r \in R} x_r(t) c_r(x_r(t))^p,$$

and so (2.3) is equivalent to showing

$$\sum_{r \in R} x_r(s) c_r(x_r(s))^p \leq \alpha^p \sum_{r \in R} x_r(s^*) c_r(x_r(s^*))^p. \tag{2.5}$$

Moreover, for $p \in (0, \infty)$ it holds that (2.1) is true if and only if

$$c_{s_i}(x_{s_i})^p \leq c_t(x_t + w_i)^p, \tag{2.6}$$

---

[8] This is the critical point where we use the fact that we consider parallel link routing games. This reasoning brakes down if strategies consist of multiple links, as then one can not write the cost objective as an additive aggregation of link contributions.

for every $i \in N$ and $t \in \mathcal{S}_i$, as the mapping $x \mapsto x^p$ is increasing in $x \geq 0$ whenever $p > 0$. This means that in order to bound the p-price of anarchy, it suffices to study the utilitarian case for the game Γ^p in which for every $r \in R$, the cost function c_r is replaced by c_r^p defined by $c_r^p(x) = c_r(x)^p$ for all $x \geq 0$. In particular, the above discussion implies that for atomic parallel link routing games

$$p\text{-PoA}(\Gamma) = \left(1\text{-PoA}(\Gamma^p)\right)^{\frac{1}{p}}. \tag{2.7}$$

A simple argument shows that upper bounds on the p-PoA for $p \in \{0, \infty\}$ can be obtained as limits of upper bounds on the p-PoA for $p \in (0, \infty)$. We will sketch this argument for $p = 0$. Let $f(p)$ be an upper bound on p-PoA(Γ). Let $s \in \text{PNE}(\Gamma)$ and s^* be a worst-case Nash equilibrium and social optimum, respectively, with respect to $C_w^0(\cdot)$. Note that s is a Nash equilibrium, independent of the value of p, and that $C_w^p(s')$ is a continuous function w.r.t. p for any fixed strategy profile s'. This means that

$$0\text{-PoA}(\Gamma) = \frac{C_w^0(s)}{C_w^0(s^*)} = \lim_{p \to 0} \frac{C_w^p(s)}{C_w^p(s^*)} \leq \lim_{p \to 0} p\text{-PoA}(\Gamma) \leq \lim_{p \to 0} f(p). \tag{2.8}$$

The first inequality follows from the fact that $C_w^p(s)$ can always be upper bounded by $C_w^p(s(p))$, where $s(p)$ is a worst-case Nash equilibrium w.r.t C_w^p for a given p, and $C_w^p(s^*)$ can be lower bounded by $C_w^p(s^*(p))$ where $s^*(p)$ is a social optimum w.r.t C_w^p for a given p. The same argument works for $p = \infty$.

2.2 Non-atomic parallel link routing game

In a non-atomic parallel link routing game there is a continuum of users that all control an infinitely small amount of flow that has to be assigned to one of the links in their strategy set. More formally, a *non-atomic parallel link routing game* $\Gamma = (R, (d_j)_{j \in [k]}, (\mathcal{S}_j)_{j \in [k]}, (c_r)_{r \in R})$ is given by a set of links R and a number of commodities $[k] = \{1, \ldots, k\}$ where a commodity has demand $d_j > 0$ for $j \in [k]$ and strategy set $\mathcal{S}_j \subseteq R$. Every link is equipped with a cost function $c_r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. A *feasible flow for commodity j* is a mapping $f^j : R \to \mathbb{R}_{\geq 0}$, where we write $f^j(r) = f_r^j$, such that $\sum_{r \in R} f_r^j = d_j$. A *feasible flow* for the non-atomic parallel link routing game is defined by the tuple $f = (f^1, \ldots, f^k)$. The *aggregation* of the feasible flows $f^1, \ldots, f^k$ for the individual commodities is defined by $f_r = f(r) = \sum_{j \in [k]} f_r^j$ for $r \in R$.

A feasible flow f is a *Wardrop flow* if for every $j \in [k]$ and $r \in \mathcal{S}_j$ with $f_r^j > 0$

$$c_r(f_r) \leq c_t(f_t) \tag{2.9}$$

for all $t \in \mathcal{S}_j$. The cost objective $C^p(f)$, the natural equivalent of the expressions in (2.10), is given by

$$C^p(f) = \begin{cases} \left(\sum_{i \in N} f_r \cdot \ell_r(f_r)^p \right)^{\frac{1}{p}} & \text{if } p \in (0, \infty) \\ \prod_{i \in N} \ell_r(f_r)^{f_r} & \text{if } p = 0 \\ \max_{i \in N} \ell_r(f_r) & \text{if } p = \infty \end{cases} \tag{2.10}$$

Similarly as for the atomic case, one may argue that in order to study the p-PoA for $p \in (0, \infty)$ it suffices to study the utilitarian case in the game Γ where all the cost functions c_r are replaced by c_r^p (and the cases $p \in \{0, \infty\}$ can be studied by taking limits). That is, we have

$$p\text{-PoA}(\Gamma) = \left(1\text{-PoA}(\Gamma^p) \right)^{\frac{1}{p}}. \tag{2.11}$$

2.3 Polynomial cost functions

In this work we will mostly be concerned with polynomial cost functions c_r of maximum degree $d \in \mathbb{N}$, i.e.,

$$c_r(x) = \sum_{j=0}^{d} a_{j,r} x^j$$

where $a_{j,r} \geq 0$ for all $j = 0, \dots, d$. We write $\mathcal{P}_d$ for the collection of all polynomials of degree at most d. We emphasize that, in general, c_r^p is not a polynomial function for $p > 0$.

3 Price of anarchy bounds

In this section we give all our unifying price of anarchy upper bounds. In Sects. 3.2 and 3.3, we rely on modified versions of lemmas from Aland et al. (2011). These are given in Appendix 2.

3.1 Non-atomic case

It is well-known that in non-atomic parallel link routing games (and more general non-atomic congestion games), the price of anarchy can be upper bounded by a so-called *smoothness parameter* defined based on the cost functions of the game. Let $\mathcal{D}$ be a class of continuous, non-decreasing cost functions. Define

$$\rho(\mathcal{D}) := (1 - \beta(\mathcal{D}))^{-1}, \quad \text{where} \quad \beta(\mathcal{D}) = \sup_{d \in \mathcal{D}} \sup_{x,y \in \mathbb{R} : x \geq y > 0} \frac{y(d(x) - d(y))}{xd(x)}. \tag{3.1}$$

Theorem 3.1 (Roughgarden [2003](#); Correa et al. [2004](#)) *Let Γ be a non-atomic parallel link routing game with cost functions from class $\mathcal{D}$. Then 1-PoA$(\Gamma) \leq \rho(\mathcal{D})$.*

The value of $\rho(\mathcal{D})$ is well-understood for many important classes of functions. Relevant to us is the class

$$\mathcal{D}_k = \{g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \ : \ g(\mu x) \geq \mu^k g(x) \ \forall \mu \in [0, 1]\}$$

for given $k \in \mathbb{R}_{\geq 0}$. In Correa et al. ([2004](#)), it is shown that

$$\beta(\mathcal{D}_k) \leq \frac{k}{(k+1)^{(k+1)/k}}$$

and thus

$$\rho(\mathcal{D}_k) \leq \left(1 - \frac{k}{(k+1)^{(k+1)/k}}\right)^{-1}.$$

An obvious fact is that for $p \in (0, \infty)$

$$c_r \in \mathcal{D}_k \Rightarrow c_r^p \in \mathcal{D}_{k \cdot p}. \tag{3.2}$$

This readily implies the following result based on (2.11).

Theorem 3.2 *Let Γ be a non-atomic parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$. Then for $p \in (0, \infty)$, we have*

$$p\text{-PoA}(\Gamma) \leq \left(1 - \frac{dp}{(dp+1)^{(dp+1)/(dp)}}\right)^{-\frac{1}{p}}. \tag{3.3}$$

An example illustrating tightness of the bound in (3.3) is given in Appendix 4. As a direct corollary, we obtain an alternative proof of the price of anarchy bound for the geometric mean objective of Vinci et al. (2022).

Corollary 3.3 (Vinci et al. [2022](#)) *Let Γ be a non-atomic parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$. The price of anarchy for the geometric mean objective satisfies*

$$0\text{-PoA}(\Gamma) \leq \left(e^{\frac{1}{e}}\right)^d \approx 1.44^d.$$

Proof Take the limit $p \to 0$ in the right hand side of (3.3). An explanation of the calculation of this limit is deferred to Proposition 5.1 in Appendix 1. $\square$

Corollary 3.4 *Let Γ be a non-atomic parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$. The price of anarchy for the egalitarian cost objective satisfies*

$$\infty\text{-PoA}(\Gamma) = 1.$$

Proof Take the limit $p \to \infty$ in the right hand side of (3.3). An explanation of the calculation of this limit is deferred to Proposition 5.1 in Appendix 1. □

Remark 3.5 (*Symmetric unweighted atomic parallel link routing games*) Fotakis (2010) showed that for symmetric unweighted atomic parallel link routing games with utilitarian cost objective, the price of anarchy can also be upper bounded by the smoothness parameter $\rho(\mathcal{D})$. This implies that the results in Theorem 3.2, Corollary 3.3 and Corollary 3.4 also hold for symmetric unweighted atomic parallel link routing games Γ.

3.2 Weighted atomic case

In order to upper bound the price of anarchy for weighted atomic parallel link routing games, we follow the ideas of Aland et al. (2011) who study (un)weighted atomic congestion games under the utalitarian cost objective. We generalize their bounds, obtained by setting $p = 1$, for the case of parallel link routing games. The main result is as follows.

Theorem 3.6 *Let Γ be a atomic weighted parallel link routing game whose cost functions are polynomials of degree at most $d \in \mathbb{N}$, and let $p \in (0, \infty)$. Let Φ_{pd} be the unique non-negative solution to $(x + 1)^{pd} = x^{pd+1}$. Then*

$$p\text{-PoA}(\Gamma) \leq \left(\Phi_{pd}^{pd+1} \right)^{\frac{1}{p}} = (\Phi_{pd} + 1)^d. \tag{3.4}$$

As a limiting case, we obtain an alternative proof of the result of Vinci et al. (2022) for the geometric mean objective.

Corollary 3.7 (Vinci et al. 2022) *Let Γ be a weighted atomic parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$. The price of anarchy for the geometric mean objective satisfies*

$$0\text{-PoA}(\Gamma) \leq 2^d.$$

Proof Rewriting the definition of Φ_{pd}, we have

$$\left(1 + \frac{1}{\Phi_{pd}} \right)^{pd} = \Phi_{pd}.$$

From this it follows directly that $\Phi_{pd} \geq 1$ for all $pd \in \mathbb{R}_{>0}$, and then also

$$\lim_{p \to 0} \Phi_{pd} = 1 \tag{3.5}$$

since $1 \leq 1 + 1/\Phi_{pd} \leq 2$. Now, by definition of Φ_{pd} it holds that $\Phi_{pd}^{pd+1} = (\Phi_{pd} + 1)^{pd}$. We then have

$$\lim_{p \to 0} \left(\Phi_{pd}^{pd+1} \right)^{\frac{1}{p}} = \lim_{p \to 0} \left((\Phi_{pd} + 1)^{pd} \right)^{\frac{1}{p}} = \lim_{p \to 0} (\Phi_{pd} + 1)^d = 2^d$$

where we use (3.5) in the last step (and the fact that $y \mapsto y^d$ is a continuous function). $\qquad \square$

We continue with the proof of Theorem 3.6.

Proof of Theorem 3.6 We start with a technical lemma to simplify the analysis. It essentially states (as will follow from the arguments that we give later) that it suffices to focus on cost functions of the form $c_r(x) = x^d$ in the main technical argument needed to prove (3.4). The proof of Lemma 3.8 uses Descartes' rule of signs (see Appendix 3), which is of independent interest. $\qquad \square$

Lemma 3.8 *Let $d \in \mathbb{N}$, $p \in (0, \infty)$, and $(\lambda, \mu) \in \mathbb{R}_{\geq 0} \times (0, 1)$ be a given pair. Then*

$$y \cdot (x + y)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \qquad (3.6)$$

for all $x, y \in \mathbb{R}_{\geq 0}$ if and only if for all $\ell \in \mathcal{P}_d$

$$y \cdot [\ell(x + y)]^p \leq \lambda \cdot y \cdot [\ell(y)]^p + \mu \cdot x \cdot [\ell(x)]^p \qquad (3.7)$$

for all $x, y \in \mathbb{R}_{\geq 0}$.

Proof First note that if (3.7) holds for all $\ell \in \mathcal{P}_d$, then (3.6) certainly also holds, by taking $\ell(x) = x^d$. We continue with the other, more involved, implication.

The statement is clearly true when $y = 0$, so assume $y > 0$. We write

$$\ell(z) = \sum_{j=0}^{d} a_j z^j$$

where $a_0, \ldots, a_d \geq 0$. Let $h := \max\{x, y\} > 0$ and set $\tilde{\ell}(z) = \tilde{a} \cdot z^d$ where

$$\tilde{a} = \frac{1}{h^d} \sum_{j=0}^{d} a_j h^j,$$

i.e., $\tilde{a}$ is chosen so that $\ell(h) = \tilde{\ell}(h)$. The polynomial

$$g(z) = \ell(z) - \tilde{\ell}(z) = (a_d - \tilde{a})x^d + \sum_{j=0}^{d-1} a_j z^j$$

has non-negative coefficients, except for $a_d - \tilde{a}$. Descartes' rule of signs (see Appendix 3) implies that g has at most one positive root (which is h). As $\tilde{\ell}(0) \leq \ell(0)$, it follows that

$$\begin{cases} \tilde{\ell}(z) \geq \ell(z) & \text{for } z \geq h \\ \tilde{\ell}(z) \leq \ell(z) & \text{for } 0 \leq z \leq h \end{cases}$$

and then also

$$\begin{cases} \tilde{\ell}(z)^p \geq \ell(z)^p & \text{for } z \geq h \\ \tilde{\ell}(z)^p \leq \ell(z)^p & \text{for } 0 \leq z \leq h \end{cases} \tag{3.8}$$

as for $p \in (0, \infty)$ the mapping $z \mapsto z^p$ is an increasing function for $z \geq 0$. Multiplying (3.6) with $\tilde{a}^p$ we obtain

$$y \cdot [\tilde{\ell}(x + y)]^p \leq \lambda \cdot y \cdot [\tilde{\ell}(y)]^p + \mu \cdot x \cdot [\tilde{\ell}(x)]^p. \tag{3.9}$$

It then follows that

$$y \cdot [\ell(x + y)]^p \leq y \cdot [\tilde{\ell}(x + y)]^p \leq \lambda \cdot y \cdot [\tilde{\ell}(y)]^p + \mu \cdot x \cdot [\tilde{\ell}(x)]^p \leq \lambda \cdot y \cdot [\ell(y)]^p$$
$$+ \mu \cdot x \cdot [\ell(x)]^p, \tag{3.10}$$

where we use (3.9) in the second inequality, and (3.8) in the first and last inequality based on the fact that $x, y \leq h \leq x + y$. This proves the claim. $\qquad\square$

Now, let $(x_r)_{r \in R}$ be the (total) flow vector of a pure Nash equilibrium s and $(y_r)_{r \in R}$ that of a system optimum s^*. Assume that for a given pair $(\lambda, \mu) \in \mathbb{R}_{\geq 0} \times (0, 1)$ it holds that

$$y_r \cdot c_r^p(x_r + y_r) \leq \lambda \cdot y_r c_r^p(y_r) + \mu \cdot x_r c_r^p(x_r) \tag{3.11}$$

for all $r \in R$. Following a standard argument, see, e.g., Aland et al. (2011, Theorem 3.2), we first apply the Nash inequality (2.1) once for every user $i \in N$, with the deviation of every user being the resource they use in the system optimum. We then find

$$c_{s_i}(x_{s_i})^p \leq c_{s_i^*}(x_{s_i^*} + w_i)^p \leq c_{s_i^*}(x_{s_i^*} + y_{s_i^*})^p$$

for every user i, where the second inequality holds, because user i uses resource s_i^* in the system optimum, and hence $w_i \leq y_{s_i^*}$. Summing up these inequalities (and rewriting the resulting sum in terms of resource loads), it then follows that

$$\sum_{r \in R} x_r c_r^p(x_r) \leq \sum_{r \in R} y_r \cdot c_r^p(x_r + y_r)$$
$$\leq \lambda \sum_{r \in R} y_r \cdot c_r^p(y_r) + \mu \sum_{r \in R} x_r \cdot c_r^p(x_r) \quad \text{(using 3.11)}$$

and, hence, by rearranging,

$$p\text{-PoA}(\Gamma) \leq \left(\frac{\lambda}{1 - \mu} \right)^{\frac{1}{p}}.$$

In order to find the best bound on the p-PoA, we consider the infimum over all possible choices of (λ, μ). By Lemma 3.8, it follows that

$$
\inf_{(\lambda,\mu)\in\mathbb{R}_{\geq 0}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall x,y\in\mathbb{R}_{\geq 0}, \ell\in\mathcal{P}_d{:}y\cdot[\ell(x+y)]^p \leq \lambda\cdot y\cdot[\ell(y)]^p + \mu\cdot x\cdot[\ell(x)]^p \right\}
$$

$$
= \inf_{(\lambda,\mu)\in\mathbb{R}_{\geq 0}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall x,y\in\mathbb{R}_{\geq 0}{:}y\cdot(x+y)^{pd} \leq \lambda\cdot y^{pd+1} + \mu\cdot x^{pd+1} \right\}
$$

$$
= \inf_{\mu\in(0,1)} \left\{ \max_{x\in\mathbb{R}_{\geq 0}} \left\{ \frac{(x+1)^{pd} - \mu\cdot x^{pd+1}}{1-\mu} \right\} \right\}
$$

$$(3.12)$$

where the last equality holds because of Lemma 6.2 in Appendix 2.

In order to analyse the infimum in (3.12), we may use Lemma 6.3 in Appendix 2. In particular, the infimum in (3.12) is equal to Φ_{pd}^{pd+1}. This then completes the proof of Theorem 3.6.

3.3 Unweighted atomic case

We continue with the price of anarchy for unweighted atomic parallel link routing games. The main result is as follows.

Theorem 3.9 *Let Γ be an atomic unweighted parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$, and let $p \in (0, \infty)$. Let Φ_{pd} be the unique non-negative solution to $(x+1)^{pd} = x^{pd+1}$ and $k = \lfloor \Phi_{pd} \rfloor$. Then*

$$
p\text{-PoA}(\Gamma) \leq \left(\frac{(k+1)^{2pd+1} - k^{pd+1}(k+2)^{pd}}{(k+1)^{pd+1} - (k+2)^{pd} + (k+1)^{pd} - k^{pd+1}} \right)^{\frac{1}{p}}. \tag{3.13}
$$

As a limiting case, we obtain the result of Vinci et al. (2022) for the geometric mean objective.

Corollary 3.10 (Vinci et al. 2022) *Let Γ be an unweighted atomic parallel link routing game whose cost functions are polynomials of degree at most d for some $d \in \mathbb{N}$. The price of anarchy for the geometric mean objective satisfies*

$$
0\text{-PoA}(\Gamma) \leq 2^d.
$$

Proof Take the limit $p \to 0$ in the right hand side of (3.13). A sketch of the calculation of this limit is deferred to Proposition 5.2 in Appendix 1. □

It is known that for the egalitarian cost objective, the price of anarchy is unbounded, already for unweighted atomic parallel link routing games with linear cost functions (Czumaj and Vöcking 2007). Indeed, the bound from Theorem 3.9 grows indefinitely as $p \to \infty$, already for $d = 1$.

Proof of Theorem 3.9 The proof uses the same approach as for the weighted case, with some differences. It is again based on the analysis of Aland et al. (2011). We start with a technical lemma that will simplify the analysis later on. It is the analogue of Lemma 3.8. Although we require $\lambda \geq 0$ in the statement of Lemma 3.11, we will in fact only apply it for the case $\lambda \geq 1$ later on. □

Lemma 3.11 *Let $d \in \mathbb{N}$, $p \in (0, \infty)$, and $(\lambda, \mu) \in \mathbb{R}_{\geq 0} \times (0, 1)$ be a given pair. Then*

$$y \cdot (x + 1)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \tag{3.14}$$

for all $x, y \in \mathbb{N}$ with $x \geq y > 0$ if and only if for all $\ell \in \mathcal{P}_d$.

$$y \cdot [\ell(x + 1)]^p \leq \lambda \cdot y \cdot [\ell(y)]^p + \mu \cdot x \cdot [\ell(x)]^p \tag{3.15}$$

for all $x, y \in \mathbb{N}$ with $x \geq y > 0$.

Proof The proof is similar to the proof of Lemma 3.8. Rouhgly speaking, the only difference is that $\}\}x + y\varepsilon$ is replaced by $\}\}x + 1\varepsilon$ in the left hand side of (3.6) and (3.7). This is not a problem as the argument in (3.10) still works, since $x, y \leq \max\{x, y\} \leq x + 1$ by assumption of $x \geq y$. □

Now, let $(x_r)_{r \in R}$ be the (total) flow vector of a pure Nash equilibrium s and $(y_r)_{r \in R}$ that of a system optimum s^*. We write $R^* = \{r \in R : y_r > 0\}$ and partition it as

$$R^+ = \{r \in R : x_r \geq y_r > 0\} \quad \text{and} \quad R^- = \{r \in R : x_r < y_r \text{ with } y_r > 0\}.$$

Assume that for a given pair $(\lambda, \mu) \in [1, \infty) \times (0, 1)$ it holds that

$$y_r \cdot c_r^p(x_r + 1) \leq \lambda \cdot y_r c_r^p(y_r) + \mu \cdot x_r c_r^p(x_r) \tag{3.16}$$

for all $r \in R^+$. Note that, as we require $\lambda \geq 1$, it also holds that

$$y_r \cdot c_r^p(x_r + 1) \leq \lambda \cdot y_r c_r^p(y_r) + \mu \cdot x_r c_r^p(x_r) \tag{3.17}$$

for all $r \in R^-$ as in that case $x_r + 1 \leq y_r$.

Again, using the same standard argument as in the proof of Theorem 3.6, we find

$$\begin{aligned}
\sum_{r \in R} x_r c_r^p(x_r) &\leq \sum_{r \in R} y_r c_r^p(x_r + 1) = \sum_{r \in R^*} y_r c_r^p(x_r + 1) \\
&= \sum_{r \in R^+} y_r c_r^p(x_r + 1) + \sum_{r \in R^-} y_r c_r^p(x_r + 1) \\
&\leq \lambda \sum_{r \in R^*} y_r c_r^p(y_r) + \mu \sum_{r \in R^*} x_r c_r^p(x_r) \\
&\leq \lambda \sum_{r \in R} y_r c_r^p(y_r) + \mu \sum_{r \in R} x_r c_r^p(x_r)
\end{aligned}$$

using (3.16) and (3.17) in the second inequality, and the fact that all cost functions are non-negative.

Rewriting gives that

$$p\text{-PoA}(\Gamma) \leq \left(\frac{\lambda}{1-\mu}\right)^{\frac{1}{p}}.$$

In order to find the best pair (λ, μ), we consider (an upper bound on) the infimum over all possible choices. By Lemma 3.11 we have that

$$\inf_{(\lambda,\mu)\in\mathbb{R}_{\geq 1}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall x, y \in \mathbb{N} \text{ with } x \geq y, \ell \in \mathcal{P}_d : y \cdot [\ell(x+1)]^p \right.$$
$$\left. \leq \lambda \cdot y \cdot [\ell(y)]^p + \mu \cdot x \cdot [\ell(x)]^p \right\}$$
$$= \inf_{(\lambda,\mu)\in\mathbb{R}_{\geq 1}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall x, y \in \mathbb{N} \text{ with } x \geq y : \; y \cdot (x+1)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \right\}$$
$$\leq \inf_{(\lambda,\mu)\in\mathbb{R}_{\geq 0}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall x, y \in \mathbb{N}_0 : \; y \cdot (x+1)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \right\}$$

$$\tag{3.18}$$

$$= \inf_{\mu\in(0,1)} \left\{ \max_{x\in\mathbb{N}_0} \left\{ \frac{(x+1)^{pd} - \mu \cdot x^{pd+1}}{1-\mu} \right\} \right\} \tag{3.19}$$

where the last equality holds because of Lemma 6.4 in Appendix 2. The inequality in (3.18), where we switch from $\lambda \geq 0$ to $\lambda \geq 1$ is justified by setting $x = 0$, as it then follows that the optimal choice of λ in (3.18) satisfies $\lambda \geq 1$.

In order to compute the final infimum in (3.19) we may use the same analysis as in the proof of Aland et al. (2011, Lemma 5.9) which also works for real values of $t = pd$ instead of only integral values of t. This computation results in the expression in the right hand side of (3.13) without the exponent $1/p$. This completes the proof of Theorem 3.9.

4 Conclusion

In this work we have given bounds on the p-price of anarchy for atomic and non-atomic parallel link routing games that unify many results in the literature for different cost objectives. To the best of our knowledge, these are the first price of anarchy bounds that unify various different well-known cost objectives such as the egalitarian, utilitarian and geometric mean objective. Our unification through the generalized p-mean objective also allows one to obtain price of anarchy bounds that can be seen as an interpolation between, e.g., the utilitarian objective and the geometric mean objective (by choosing $p \in [0, 1]$). Understanding how the price of anarchy behaves under different cost objectives is important to decide how inefficient (from a quantitative point of view) an equilibrium flow actually is, and whether or not measures are needed to reduce it.

We note that Vinci et al. (2022) also obtain a result for the geometric mean objective in the setting of *online load balancing* using a *greedy algorithm*. In Appendix 5 we show that there exists a unification of their result and that of Caragiannis (2008),

who studies the same greedy algorithm under the L_p-norm objective, when the links $r \in R$ have cost functions of the form $c_r(x) = a_r x^d$ for a given $d \in \mathbb{N}$. We suspect a similar result holds true for general polynomial cost functions of degree at most d, but leave this as a problem for future work. Further related works on online load balancing are, e.g., Awerbuch et al. (1995), Bilò and Vinci (2017), Klimm et al. (2019).

Appendix 1: Limit calculations

Proposition 5.1 *Let $d > 0$ be given. Then*

$$\lim_{p \to 0} \left(1 - \frac{dp}{(dp+1)^{(dp+1)/(dp)}} \right)^{-\frac{1}{p}} = \left(e^{\frac{1}{e}} \right)^d$$

and

$$\lim_{p \to \infty} \left(1 - \frac{dp}{(dp+1)^{(dp+1)/(dp)}} \right)^{-\frac{1}{p}} = 1$$

Proof Setting $t = dp$, we obtain

$$\left(1 - \frac{dp}{(dp+1)^{(dp+1)/(dp)}} \right)^{-\frac{1}{p}} = \left[\left(1 - \frac{t}{(t+1)^{(t+1)/t}} \right)^{-\frac{1}{t}} \right]^d$$

It now suffices to study the limit $t \to 0$, or $t \to \infty$, as d is fixed. We have

$$\left(1 - \frac{t}{(t+1)^{(t+1)/t}} \right)^{-\frac{1}{t}} = e^{-\frac{\ln\left(1 - \frac{t}{(t+1)^{(t+1)/t}} \right)}{t}}$$

Therefore, it suffices to show that

$$\lim_{t \to 0} \frac{\ln\left(1 - \frac{t}{(t+1)^{(t+1)/t}} \right)}{t} = -\frac{1}{e} \tag{5.1}$$

and

$$\lim_{t \to \infty} \frac{\ln\left(1 - \frac{t}{(t+1)^{(t+1)/t}} \right)}{t} = 0 \tag{5.2}$$

In order to prove these identities, one may use L'Hôpital's rule (i.e., replace the functions in the nominator and denominator by their derivatives). These are somewhat

tedious, but elementary, calculus exercises for which we only give informal sketches here.

For $t \to 0$, note that $(t+1)^{\frac{t+1}{t}} = (1+t)(1+t)^{\frac{1}{t}} \approx e$ as $\lim_{t\to 0}(1+t)^{\frac{1}{t}} = e$. Substituting this in (5.1) gives

$$\lim_{t\to 0} \frac{\ln\left(1 - \frac{t}{(t+1)^{(t+1)/t}}\right)}{t} \approx \lim_{t\to 0} \frac{\ln\left(1 - \frac{t}{e}\right)}{t} = \lim_{t\to 0} \frac{-\frac{1}{e}}{1 - \frac{t}{e}} = -\frac{1}{e}$$

where the second to last equality is an application of L'Hôpital's rule.

For $t \to \infty$, note that $(t+1)^{\frac{t+1}{t}} = (1+t)(1+t)^{\frac{1}{t}} \approx 1+t$ as $\lim_{t\to\infty}(1+t)^{\frac{1}{t}} = 1$. Then (5.2) becomes

$$\lim_{t\to\infty} \frac{\ln\left(1 - \frac{t}{(t+1)^{(t+1)/t}}\right)}{t} \approx \lim_{t\to\infty} \frac{\ln\left(1 - \frac{t}{t+1}\right)}{t} = \lim_{t\to\infty} \frac{\ln\left(\frac{1}{t+1}\right)}{t} = \lim_{t\to\infty} \frac{t+1}{-(1+t)^2} = 0,$$

where again the second to last equality is an application of L'Hôpital's rule. □

Proposition 5.2 *Let $d > 0$ be given and let $p \in (0, \infty)$. Let Φ_{pd} be the unique nonnegative solution to $(x+1)^{pd} = x^{pd+1}$ and $k = \lfloor \Phi_{pd} \rfloor$. Then*

$$\lim_{p\to 0} \left(\frac{(k+1)^{2pd+1} - k^{pd+1}(k+2)^{pd}}{(k+1)^{pd+1} - (k+2)^{pd} + (k+1)^{pd} - k^{pd+1}} \right)^{\frac{1}{p}} = 2^d$$

Proof Setting $t = dp$, we obtain

$$\left(\frac{(k+1)^{2pd+1} - k^{pd+1}(k+2)^{pd}}{(k+1)^{pd+1} - (k+2)^{pd} + (k+1)^{pd} - k^{pd+1}} \right)^{\frac{1}{p}}$$

$$= \left[\left(\frac{(k+1)^{2t+1} - k^{t+1}(k+2)^{t}}{(k+1)^{t+1} - (k+2)^{t} + (k+1)^{t} - k^{t+1}} \right)^{\frac{1}{t}} \right]^{d}.$$

Therefore, it suffices to show

$$\lim_{t\to 0} \left(\frac{(k+1)^{2t+1} - k^{t+1}(k+2)^{t}}{(k+1)^{t+1} - (k+2)^{t} + (k+1)^{t} - k^{t+1}} \right)^{\frac{1}{t}} = 2 \qquad (5.3)$$

For $t = pd$ small enough, it holds that $k = \lfloor \Phi_{pd} \rfloor = 1$. To see this, note that if $x \geq 0$ is such that $(x+1)^t = x^{t+1}$ then we must have $x \geq 1$, and $x = x(t)$ is increasing in t. Then (5.3) reduces to

$$\lim_{t \to 0} \left(\frac{2^{2t+1} - 3^t}{2^{t+1} - 3^t + 2^t - 1} \right)^{\frac{1}{t}} = 2 \tag{5.4}$$

This identity can be shown using L'Hôpital's rule after exponentiating the left hand side similar to what was done in the proof of Proposition 5.1. That is, it suffices to show that

$$\lim_{t \to 0} \frac{\ln \left(\frac{2^{2t+1} - 3^t}{2^{t+1} - 3^t + 2^t - 1} \right)}{t} = \lim_{t \to 0} \frac{\ln \left(2^{2t+1} - 3^t \right)}{t} - \lim_{t \to 0} \frac{\ln \left(2^{t+1} - 3^t + 2^t - 1 \right)}{t} = \ln(2)$$

Using L'Hôpital's rule, we have

$$\lim_{t \to 0} \frac{\ln \left(2^{2t+1} - 3^t \right)}{t} = \lim_{t \to 0} \frac{2 \ln(4)4^t - \ln(3)3^t}{2^{2t+1} - 3^t} = 2\ln(4) - \ln(3) = \ln(16/3)$$

and

$$\lim_{t \to 0} \frac{\ln \left(2^{t+1} - 3^t + 2^t - 1 \right)}{t} = \lim_{t \to 0} \frac{2 \ln(2)2^t - \ln(3)3^t + \ln(2)2^t - 1}{2^{t+1} - 3^t + 2^t - 1} = 2\ln(2) - \ln(3) + \ln(2)$$

Substracting $2\ln(2) - \ln(3) + \ln(2) = \ln(8/3)$ from $\ln(16/3)$ then indeed gives us $\ln(2)$, as desired. □

Appendix 2: Lemmas from Aland et al. (2011)

This section contains all the lemmata from Aland et al. (2011) that we rely on. Lemma 6.1 can be proved using the proof of Aland et al. (2011, Lemma 5.2) verbatim.

Lemma 6.1 *Let $\mu \in (0, 1]$ and $t > 0$. Define $g : \mathbb{R}_{\geq 0} \to \mathbb{R}$ as $g(x) = (x + 1)^t - \mu \cdot x^{t+1}$. Then it holds that g has exactly one local maximum at some $\xi \in \mathbb{R}_{>0}$. Moreover, g is strictly increasing in $[0, \xi)$, and strictly decreasing in (ξ, ∞).*

The following lemma is a relaxed version of Aland et al. (2011, Lemma 5.5).

Lemma 6.2 *Let $t = pd > 0$. Then*

$$\inf_{(\lambda, \mu) \in \mathbb{R}_{\geq 0} \times (0,1)} \left\{ \frac{\lambda}{1 - \mu} \,\middle|\, \forall x, y \in \mathbb{R}_{\geq 0} : y \cdot (x + y)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \right\}$$

$$= \inf_{\mu \in (0,1)} \left\{ \max_{x \in \mathbb{R}_{\geq 0}} \left\{ \frac{(x + 1)^{pd} - \mu \cdot x^{pd+1}}{1 - \mu} \right\} \right\}$$

$$\tag{6.1}$$

Fig. 3 Parallel link network consisting of two arcs

$$c_T(x) = x^d$$

$$c_B(x) = 1$$
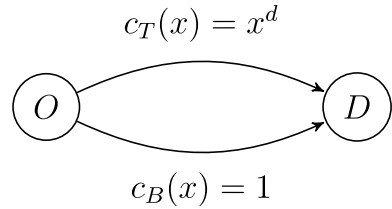
In order to prove Lemma 6.2 one may follow the proof of Aland et al. (2011, Lemma 5.5). Aland et al. (2011) give the proof for $t \in \mathbb{N}$, but the only time this assumption is used, is when the proof relies on Aland et al. (2011, Lemma 5.5), which in fact also holds for any real $t > 0$, as summarized in Lemma 6.1. Furthermore, the statement of Aland et al. (2011, Lemma 5.5) is phrased in terms of general polynomials, but the proof proceeds by making a reduction to monomials with coefficient 1, which is our starting point (as we reduced to monomials already in Lemma 3.8).

Lemma 6.3 *Let* $t = pd > 0$. *Then*

$$\inf_{\mu \in (0,1)} \left\{ \max_{x \in \mathbb{R}_{\geq 0}} \left\{ \frac{(x+1)^{pd} - \mu \cdot x^{pd+1}}{1 - \mu} \right\} \right\} = \Phi_{pd}^{pd+1}.$$

In order to prove Lemma 6.3 we may use the calculus proof of Aland et al. (2011, Lemma 5.6) by observing that it also verbatim works for real values of $t = pd$ instead of only $t \in \mathbb{N}$ by using Lemmas 6.1 and 6.2 in the argumentation as opposed to Aland et al. (2011, Lemmas 5.2 and 5.5), respectively.

Lemma 6.4 *Let* $t = pd > 0$. *Then*

$$\inf_{(\lambda,\mu) \in \mathbb{R}_{\geq 0} \times (0,1)} \left\{ \frac{\lambda}{1 - \mu} \; \middle| \; \forall x, y \in \mathbb{N}_0 : \; y \cdot (x+1)^{pd} \leq \lambda \cdot y^{pd+1} + \mu \cdot x^{pd+1} \right\}$$

$$= \inf_{\mu \in (0,1)} \left\{ \max_{x \in \mathbb{N}_0} \left\{ \frac{(x+1)^{pd} - \mu \cdot x^{pd+1}}{1 - \mu} \right\} \right\}$$

$$= \frac{(k+1)^{2pd+1} - k^{pd+1}(k+2)^{pd}}{(k+1)^{pd+1} - (k+2)^{pd} + (k+1)^{pd} - k^{pd+1}}.$$

The first equality can be proved in a similar way as the proofs of Aland et al. (2011, Lemmas 5.7 and 5.8). The second equality can be proved similarly as Aland et al. (2011, Lemmas 5.9). The proofs of these lemmas also hold verbatim for real $t = pd$ instead of only $t \in \mathbb{N}$.

Appendix 3: Descartes' rule of signs

Descartes' rule of signs can be used to upper bound the number of positive roots of a polynomial in terms of the number of sign changes of the coefficients of the polynomial. The version that we use of it in this work can be summarized as follows.

Let $f(x) = a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0$ be a polynomial of degree (at most) n with real coefficients $a_n, \ldots, a_0$. Define

$$SC(f) = \{i \in [n] : \mathrm{sgn}(a_i) \neq \mathrm{sgn}(a_{i-1})\}$$

where $[n] = \{1, 2, \ldots, n\}$. Here we use the sign function $\mathrm{sgn} : \mathbb{R} \to \{-1, +1\}$ defined by $\mathrm{sgn}(a) = 1$ if $a > 0$ and $\mathrm{sgn}(a) = -1$ if $a < 0$. We refer to $|SC(f)|$ as the number of *sign changes* of f. Descartes' rule of signs (in particular) says that the number of positive real roots of f is at most the number of sign changes of f.

Appendix 4: Lower bound for non-atomic games

In this section we give a lower bound on a parallel link network with two links, illustrating the tightness of the bound in Theorem 3.2. This is the same example as used in Correa et al. (2004); Roughgarden (2003). We have one link B with $c_B(x) = 1$ and one link T with $c_T(x) = x^d$; see also Fig. 3. There is one commodity with a total flow of $d_1 = 1$. It is not hard to see that a Wardrop flow f is given by routing the whole unit of flow of the commodity over T. The social optimum w.r.t. $C^p(f)$ is given by routing $1 - dp(dp+1)^{\frac{dp+1}{dp}}$ units of flow over link T, and the remainder over link B. A simple calculation of the resulting price of anarchy shows that this gives the desired bound as in Theorem 3.2.

Appendix 5: Online load balancing and the greedy algorithm

An instance of *online load balancing* consists of the same components as an atomic weighted parallel link routing game, but proceeds differently. The users in $N = \{1, \ldots, n\}$, often called *clients*, arrive online and have to irrevocably be assigned to some link $r \in R$ from their strategy set. We assume without loss of generality that the clients arrive according to the permutation $(1, 2, \ldots, n)$, so in step t of this process, client t has to be assigned. The goal is to find an assignment that minimizes the social cost, which in our case, based on (2.7), is equivalent to

$$\left(\sum_{r \in R} x_r c_r^p(x_r) \right)^{\frac{1}{p}}$$

The greedy algorithm, as the name suggests, makes the assignment decision "greedily" meaning that we assign a client to a link in her strategy set that yields the smallest increase in cost objective C_w^p (based on the clients that have been assigned so far).

To be precise, suppose that $s^{(t)} = (s_1, \ldots, s_t) \in \times_{i=1}^t \mathcal{S}_i$ is the profile in which clients $1, \ldots, t$ have been assigned so far. Then we assign client $t + 1$ to the link

$$s_{t+1} = \text{argmin}_{r \in \mathcal{S}_{t+1}} C_w^p(s^{(t)}, r) \tag{5.1}$$

If $p \in (0, \infty)$, we equivalently have

$$s_{t+1} = \text{argmin}_{r \in \mathcal{S}_{t+1}} \sum_{r \in R} x_r c_r^p(x_r) \tag{5.2}$$

as the function $y \mapsto y^p$ is increasing in $y \geq 0$.

If, for some $d \in \mathbb{N}$, we have $c_r(x) = a_r x^d$ with $a_r > 0$ for every $r \in R$, then we have

$$\sum_{r \in R} x_r c_r^p(x_r) = \sum_{r \in R} a_r x_r^{dp+1} = \sum_{r \in R} (a_r^{1/(dp+1)} x_r)^{dp+1}$$

We define

$$w_{ir} = a_r^{1/(dp+1)} w_i$$

for $r \in R$. With the definition of w_{ir} in hand, we can rephrase the social cost as

$$\left[\left(\sum_{r \in R} \ell_r(z_r(s))^{1+dp} \right)^{\frac{1}{1+dp}} \right]^{\frac{dp+1}{p}} \tag{5.3}$$

where $z_r(s) = \sum_{i:s_i=r} w_{ir}$ is the sum of the weights of the clients assigned to resource r. If we define $q = 1 + dp$, we have arrived at (a special case of) the setting of Caragiannis (2008). To be more precise, (5.3) corresponds to studying the L_{dp+1}-norm cost objective raised to the power $(dp + 1)/p$.[9]

We next restate (Caragiannis 2008, Theorem 3.1) in a slightly different form. The guarantee given in (5.4) is the same as that in Caragiannis (2008, Theorem 3.1), with the parameter p in Caragiannis (2008) replaced by $dp + 1$, and raised to the power $(dp + 1)/p$ as a consequence of (5.3).

Theorem 5.1 (Caragiannis 2008) *For $d \in \mathbb{N}$ and $p \in (0, \infty)$, consider an instance of online load balancing with cost functions of the form $c_r(x) = a_r x^d$ for $r \in R$. Let s be the strategy profile obtained with the greedy algorithm defined by the rule in (5.2), and let s^* be a system optimum. Then*

$$\frac{C_w^p(s)}{C_w^p(s^*)} \leq (2^{1/(dp+1)} - 1)^{-(dp+1)/p}. \tag{5.4}$$

---

[9] Our parameter $dp + 1$ plays the role of the parameter "p" in Caragiannis (2008).

As a corollary we obtain the bound of Vinci et al. (2022) for this class of cost functions. We note that their result for $p = 0$ holds for arbitrary polynomials of degree d with non-negative coefficients.

Corollary 5.2 *For $d \in \mathbb{N}$, consider an instance of online load balancing with cost functions of the form $c_r(x) = a_r x^d$ for $r \in R$. Let s be the strategy profile obtained with the greedy algorithm defined by the rule in (5.2), and let s^* be a system optimum. Then*

$$\frac{C_w^0(s)}{C_w^0(s^*)} \le 4^d. \tag{5.5}$$

Proof We take the limit $p \to 0$ on the right hand side of (5.4), i.e.,

$$\lim_{p \to 0}(2^{1/(dp+1)} - 1)^{-(dp+1)/p} = 4^d$$

which can be shown by setting $t = pd$, and then exponentiating and applying L'Hôpital's rule (as for the limits analysed in Appendix 1). That is, we can write

$$\left(2^{1/(dp+1)} - 1\right)^{-\frac{dp+1}{p}} = \left[\left(2^{\frac{1}{t+1}} - 1\right)^{\frac{t+1}{t}}\right]^d.$$

It suffices to show that

$$\lim_{t \to 0} \frac{t+1}{t} \ln\left(2^{\frac{1}{t+1}} - 1\right) = \ln(4)$$

We have

$$\frac{t+1}{t} \ln\left(2^{\frac{1}{t+1}} - 1\right) = \ln\left(2^{\frac{1}{t+1}} - 1\right) + \frac{\ln\left(2^{\frac{1}{t+1}} - 1\right)}{t}$$

Taking $t \to 0$ in the first term gives zero. Applying L'Hôpital's rule to the second term gives

$$\lim_{t \to 0} \frac{\ln\left(2^{\frac{1}{t+1}} - 1\right)}{t} = \lim_{t \to 0} \frac{2\ln(2)\left(\frac{1}{2}\right)^t}{2^{\frac{1}{t+1}} - 1} = 2\ln(2) = \ln(4)$$

which gives the desired result. $\qquad \square$

References

Aland S, Dumrauf D, Gairing M, Monien B, Schoppmann F (2011) Exact price of anarchy for polynomial congestion games. SIAM J Comput 40(5):1211–1233

Anshelevich E, Dasgupta A, Kleinberg J, Tardos É, Wexler T, Roughgarden T (2008) The price of stability for network design with fair cost allocation. SIAM J Comput 38(4):1602–1623

Awerbuch B, Azar Y, Grove EF, Kao M-Y, Krishnan P, Vitter JS (1995) Load balancing in the l/sub p/norm. In: Proceedings of IEEE 36th annual foundations of computer science. IEEE, pp 383–391

Bhawalkar K, Gairing M, Roughgarden T (2014) Weighted congestion games: the price of anarchy, universal worst-case examples, and tightness. ACM Trans Econ Comput 2(4):1–23

Bilò V, Vinci C (2017) On the impact of singleton strategies in congestion games. In: 25th annual European symposium on algorithms. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik

Bilò V, Vinci C (2019) Dynamic taxes for polynomial congestion games. ACM Trans Econ Comput 7(3):1–36

Bonifaci V, Salek M, Schäfer G (2011) Efficiency of restricted tolls in non-atomic network routing games. In: International symposium on algorithmic game theory. Springer, pp 302−313

Braess D, Nagurney A, Wakolbinger T (2005) On a paradox of traffic planning. Transp Sci 39(4):446–450

Caragiannis I (2008) Better bounds for online load balancing on unrelated machines. In: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms, pp 972–981

Caragiannis I, Kaklamanis C, Kanellopoulos P (2010) Taxes for linear atomic congestion games. ACM Trans Algorithms 7(1):13:1-13:31

Caragiannis I, Flammini M, Kaklamanis C, Kanellopoulos P, Moscardelli L (2011) Tight bounds for selfish and greedy load balancing. Algorithmica 61(3):606–637

Chen P-A, de Keijzer B, Kempe D, Schäfer G (2014) Altruism and its impact on the price of anarchy. ACM Trans Econ Comput 2(4):171–1745

Christodoulou G, Gairing M (2015) Price of stability in polynomial congestion games. ACM Trans Econ Comput 4(2):1–17

Christodoulou G, Koutsoupias E (2005) The price of anarchy of finite congestion games. In: Proceedings of the thirty-seventh annual ACM symposium on theory of computing, pp 67–73

Christodoulou G, Koutsoupias E, Spirakis PG (2011) On the performance of approximate equilibria in congestion games. Algorithmica 61(1):116–140

Christodoulou G, Gairing M, Giannakopoulos Y, Spirakis PG (2019) The price of stability of weighted congestion games. SIAM J Comput 48(5):1544–1582

Colini-Baldeschi R, Cominetti R, Scarsini M (2019) Price of anarchy for highly congested routing games in parallel networks. Theory Comput Syst 63(1):90–113

Cominetti R, Scarsini M, Schröder M, Stier-Moses NE (2019) Price of anarchy in stochastic atomic congestion games with affine costs. In: Proceedings of the 2019 ACM conference on economics and computation, pp 579–580

Correa JR, Schulz AS, Stier-Moses NE (2004) Selfish routing in capacitated networks. Math Oper Res 29(4):961–976

Czumaj A, Vöcking B (2007) Tight bounds for worst-case equilibria. ACM Trans Algorithms 3(1):1–17

Dubey P (1986) Inefficiency of nash equilibria. Math Oper Res 11(1):1–8

Fotakis D (2010) Congestion games with linearly independent paths: convergence time and price of anarchy. Theory Comput Syst 47(1):113–136

Fotakis D, Kalimeris D, Lianeas T (2015) Improving selfish routing for risk-averse players. In: International conference on web and internet economics. Springer, pp 328–342

Gairing M, Lücking T, Mavronicolas M, Monien B (2006) The price of anarchy for polynomial social cost. Theoret Comput Sci 369(1–3):116–135

Gairing M, Lücking T, Mavronicolas M, Monien B, Rode M (2008) Nash equilibria in discrete routing games with convex latency functions. J Comput Syst Sci 74(7):1199–1225

Gairing M, Schoppmann F (2007) Total latency in singleton congestion games. In: International workshop on web and internet economics. Springer, pp 381–387

Harks T, Schröder M, Vermeulen D (2019) Toll caps in privatized road networks. Eur J Oper Res 276(3):947–956

John F (1950) Nash. The bargaining problem. Econometrica 18(2):155–162

Kleer P, Schäfer G (2017) Path deviations outperform approximate stability in heterogeneous congestion games. In: International symposium on algorithmic game theory. Springer, pp 212–224

Kleer P, Schäfer G (2019a) The impact of worst-case deviations in non-atomic network routing games. Theory Comput Syst 63(1):54–89

Kleer P, Schäfer G (2019b) Tight inefficiency bounds for perception-parameterized affine congestion games. Theoret Comput Sci 754:65–87

Kleer P, Schäfer G (2021) Computation and efficiency of potential function minimizers of combinatorial congestion games. Math Program 190(1):523–560

Klimm M, Schmand D, Tönnis A (2019) The online best reply algorithm for resource allocation problems. In: International symposium on algorithmic game theory. Springer, pp 200–215

Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. Comput Sci Rev 3(2):65–69

Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. In: Proceedings of the 16th annual conference on theoretical aspects of computer science, pp 404–413

Lianeas T, Nikolova E, Stier-Moses NE (2019) Risk-averse selfish routing. Math Oper Res 44(1):38–57

Monnot B, Benita F, Piliouras G (2017) Routing games in the wild: efficiency, equilibration and regret. In: International conference on web and internet economics. Springer, pp 340–353

Nikolova E, Stier-Moses NE (2014) A mean-risk model for the traffic assignment problem with stochastic travel times. Oper Res 62(2):366–382

Pigou AC (1920) The economics of welfare. Macmillan, New York

Piliouras G, Nikolova E, Shamma JS (2016) Risk sensitivity of price of anarchy under uncertainty. ACM Trans Econ Comput 5(1):51–527

Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. Int J Game Theory 2:65–67

Roughgarden T (2003) The price of anarchy is independent of the network topology. J Comput Syst Sci 67(2):341–364

Roughgarden T (2005) Selfish routing and the price of anarchy. MIT Press, Cambridge

Roughgarden T (2015) Intrinsic robustness of the price of anarchy. J ACM 62(5):32

Roughgarden T, Tardos É (2002) How bad is selfish routing? J ACM 49(2):236–259

Schröder M (2020) Price of anarchy in congestion games with altruistic/spiteful players. In: International symposium on algorithmic game theory. Springer, pp 146−159

Takalloo M, Kwon C (2020) On the price of satisficing in network user equilibria. Transp Sci 54(6):1555–1570

United States. Bureau of Public Roads (1964) Traffic assignment manual for application with a large, high speed computer, vol 37. US Department of Commerce, Bureau of Public Roads (1964)

Vinci C, Bilò V, Monaco G, Moscardelli L (2022) Nash social welfare in selfish and online load balancing. ACM Trans Econ Comput 10(2):41. https://dl.acm.org/doi/full/10.1145/3544978

Wardrop JG (1952) Some theoretical aspects of road traffic research. Proc Inst Civ Eng 1:325–378

Zhang J, Pourazarm S, Cassandras CG, Paschalidis IC (2018) The price of anarchy in transportation networks: data-driven evaluation and reduction strategies. Proc IEEE 106(4):538–553