



Data-Driven Models of Selfish Routing: Why Price of Anarchy Does Depend on Network Topology

Francisco Benita¹, Vittorio Bilò², Barnabé Monnot³,
Georgios Piliouras¹, and Cosimo Vinci⁴

¹ Singapore University of Technology and Design, Singapore, Singapore
{francisco_benita,georgios}@sutd.edu.sg

² University of Salento, Lecce, Italy
vittorio.bilo@unisalento.it


³ Ethereum Foundation, Zug, Switzerland
barnabemonnot@gmail.com

⁴ Gran Sasso Science Institute, L'Aquila, Italy
cosimo.vinci@gssi.it

Abstract. We investigate traffic routing both from the perspective of real world data as well as theory. First, we reveal through data analytics a natural but previously uncaptured regularity of real world routing behavior. Agents only consider, in their strategy sets, paths whose free-flow costs (informally their lengths) are within a small multiplicative $(1 + \theta)$ constant of the optimal free-flow cost path connecting their source and destination where $\theta \geq 0$. In the case of Singapore, $\theta = 1$ is a good estimate of agents' route (pre)selection mechanism. In contrast, in Pigou networks the ratio of the free-flow costs of the routes and thus θ is *infinite*, so although such worst case networks are mathematically simple they correspond to artificial routing scenarios with little resemblance to real world conditions, opening the possibility of proving much stronger Price of Anarchy guarantees by explicitly studying their dependency on θ . We provide an exhaustive analysis of this question by providing provably tight bounds on $\text{PoA}(\theta)$ for arbitrary classes of cost functions both in the case of general congestion/routing games as well as in the special case of path-disjoint networks. For example, in the case of the standard Bureau of Public Roads (BPR) cost model, $c_e(x) = a_e x^4 + b_e$ and more generally quartic cost functions, the standard PoA bound for $\theta = \infty$ is 2.1505 [21] and it is tight both for general networks as well as path-disjoint and even parallel-edge networks. In comparison, in the case of $\theta = 1$, the PoA in the case of general networks is only 1.6994, whereas

Full paper can be found at <https://arxiv.org/abs/2009.12871>. F. Benita would like to acknowledge Ministry of Education, Singapore Grant SGPCTRS1804. B. Monnot acknowledges the SUTD Presidential Graduate Fellowship. G. Piliouras gratefully acknowledges AcRF Tier 2 grant 2016-T2-1-170, grant PIE-SGP-AI-2020-01, NRF2019-NRF-ANR095 ALIAS grant and NRF 2018 Fellowship NRF-NRFF2018-07. C. Vinci would like to acknowledge the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Markets”.

for path-disjoint/parallel-edge networks is even smaller (1.3652), showing that both the route geometries as captured by the parameter θ as well as the network topology have significant effects on PoA (Fig. 1).



(a) Comparison between PoA(θ) in the ver- case of quartic costs for general/path- sus chosen path. Discussion on disjoint networks resp. and the standard data analytics can be found in bound $PoA(\infty)=2.1505$ from [21]. More Section 2. results can be found in Table 1.

Fig. 1. Improved Price of Anarchy bounds in data-driven routing models

1 Introduction

Modern cities are wonders of emergent, largely self-organizing, behavior. Major capitals buzz with the collective hum of millions of people whose lives are intertwined and coupled in myriad and diverse ways. One of the most palpable such phenomena of collective behavior is the emergence and diffusion of traffic throughout the city. A bird’s eye view of any major city would reveal a complex and heterogeneous landscape of thousands upon thousands of cars, buses, trucks, motorcycles, running though the veins of a maze of remarkable complexity and scale consisting of a vast number of streets and highways. As Fig. 2 suggests, the full magnitude of the multi-scale complexity of these real-life networks lies outside the perceptive capabilities of any single individual. Nevertheless, as a phenomenon that we get to experience daily, such as the weather, we would like to understand at least some macroscopic, high level characteristics of traffic routing. Quite possibly, one of the most interesting such questions is how efficient is a traffic network?

This question has received a lot of attention within algorithmic game theory. Using the model of congestion games, seminal papers in the area established tight bounds on their Price of Anarchy (PoA), i.e., the worst case inefficiency of traffic routing [13, 23]. For example, the Price of Anarchy of linear non-atomic congestion games is $4/3$, whereas if we apply the standard Bureau of Public Roads (BPR) cost functions that are polynomials of degree four, then the Price of Anarchy is roughly 2.151. On the positive side, these bounds apply to all networks (within the prescribed class of delay/cost functions) regardless of their

size or their total demand, or number of agents and are tight even for the simplest possible network instances, i.e., Pigou networks with just two parallel links.

The common interpretation of these bounds is that they are strong and a PoA anywhere in that range (e.g. $\text{PoA} = 2$) immediately translates to practical guarantees about real traffic. Some recent purely experimental work, however, has produced new insights that allow us to reexamine these results from a different perspective. For example, [16] showed that the efficiency of real-life traffic networks, as estimated from traffic measurements alone, is really close to optimal even when compared to very optimistic estimates of optimal performance. A Price of Anarchy of 2 implies that the average commuter can increase their mean speed by 100%. Measurements suggest that this level of inefficiencies/improvements is rather unlikely. Since Price of Anarchy is a macroscopic characteristic of a system with countless moving parts, a more useful analogy is that of weather or climate (e.g., average temperature). The differences between 10% and 20% increase to system inefficiency are significant and a 100% increase, i.e., PoA of 2 would have catastrophic consequences.

A Natural Question Emerges: Can we create classes of models, i.e., congestion games, which come closer to representing real world traffic? In this paper we do, by leveraging an intuitive but largely unexplored characteristic of real world traffic routing. Commuters only consider in their strategy sets paths/routes whose free-flow costs (informally their lengths) are approximately equal to each other (within a multiplicative factor of $1 + \theta$). We call such games θ -free flow games. We generalize the special case of linear congestion θ -free flow games [4] to the case of arbitrary classes of cost functions as well as simultaneously studying both general networks as well as path-disjoint networks. $\theta = 0$ means that all paths considered by each user have exactly equal free-flow cost/length, whereas $\theta = 1$ allows for paths whose lengths are within a factor of 2. Pigou networks may feel intuitively very simple and thus natural due to their small size, but they fail to satisfy this property in the most extreme sense. The ratio of the free-flow costs of the two edges is infinite ($\theta = \infty$). It is like considering two possible paths from home to work, one which is the shortest distance route and one that circumnavigates the globe along the way. Such unnatural paths may indeed be available to us, but we unconsciously and automatically prune them out from the set of alternatives that we consider. Amazingly, enforcing such a natural property on the set of models (routing games) we consider immediately removes from consideration Pigou networks, the worst case examples from a PoA perspective, and thus opens up the possibility of proving stronger Price of Anarchy guarantees. What are the implications of such characteristics to PoA? What other type of attributes can we take advantage of when creating new models? Finally, how well do they match real traffic conditions?

We hope that this paper opens up a new direction for tighter coupling between data analytics, modelling and theory in congestion games and beyond. Analyzing different cities as well as introducing models that take into account the difference between public and private transport seem like an exciting direction for future work.

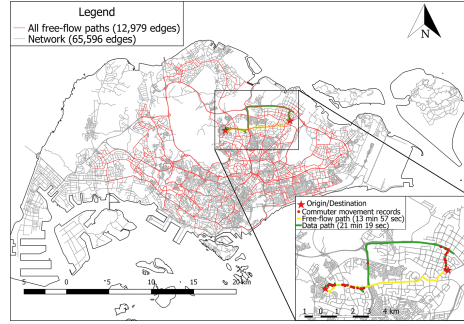


Fig. 2. For each trip segment, we find the best free-flow time and the data free-flow time. The reconstruction of the selected route uses datapoints logged along the trip. In yellow, the fastest route in free-flow condition is highlighted. The reconstructed route is in green, along which we find the data free-flow time.

1.1 Our Contribution

In Sect. 2, we start off by experimentally computing estimates of θ from real world traffic data. We employ an experimental dataset that contains detailed information (sampled every 13 s) on the routing behavior of tens of thousands of commuters in Singapore. Based on this fine-grained information and in combination with a graph representation of the road network of Singapore that we have created we can estimate numerous characteristics of the actual routing behavior at an unprecedented level of accuracy. Using these tools that we believe are of independent interest as well, we find that the θ values for the vast majority of commuters (close to 80%) are below 1.

Inspired by the above evidence, we introduce a new class of congestion games, that we call free-flow games, parametrized by θ (Sect. 3). We provide two parametric *tight bounds* on the Price of Anarchy of free-flow games under general latency functions satisfying mild assumptions, thus largely extending the results given in [4] which are restricted to affine latencies only. The first of these bounds applies to the general case of unrestricted network topologies (indeed, it applies even to congestion games) (Theorem 1), while the second one holds for path-disjoint networks (Theorem 2) which includes the fundamental parallel-links topology. These bounds are never equal as long as $\theta \notin \{0, \infty\}$. In fact, differently from what happens in the classical setting without the free-flow assumption, where the worst-case situation already arises in a two parallel-links network (the Pigou network), for free-flow games the absence of intersections among paths allows for more efficient equilibria. More precisely, as θ goes to infinity, both bounds converge to the same limit, but the convergence of the one for parallel-link networks can be significantly slower (see, for instance, Fig. 1(a)). We also stress that, with respect to the case of affine latency functions, our findings improve on the results given in [4], as we close the gap between upper and lower bound on the Price of Anarchy for parallel-link networks that was left as an open problem.

One of the most important messages coming from our investigation is that the separation outlined by Theorems 1 and 2 sheds new light on the question of whether the Price of Anarchy is affected by the network topology. In fact, a famous, and perhaps counter-intuitive, result by Roughgarden [21] states that the PoA is independent of the network topology as, in almost all notable cases, worst-case instances are already attained by simple networks, such as parallel-link graphs. Under the free-flow assumption, however, this situation ceases to hold, and the network topology begins to play a critical, if not dominant, role in the efficiency of equilibria. This evidence has major practical implications, as it signifies the fundamental importance of careful road network design and planning for selfish routing. As shown in Fig. 1 and in more details in Table 1, in the case of the standard Bureau of Public Roads (BPR) cost model, $c_e(x) = a_e x^4 + b_e$ and more generally quartic cost functions, applying the constraint $\theta = 1$ nearly halves the percentage of inefficiency, and applying the additional constraint of a path-disjoint network halves it once again.

At the technical level, our general formulas depend on whether the free-flow traversing time of some edges is larger than zero, i.e., whether the limit of the edge cost/latency as its load goes to zero is strictly positive. Latency functions for which this does not hold have been termed *homogeneous* by Roughgarden [21] and they represent one of the few exceptions for which he could not prove that the PoA is independent of the network topology. Since under homogeneous latency functions any congestion game is a 0-free flow game, as a by-product of our results, we also obtain that, for (free-flow) games with homogeneous latency functions, the Price of Anarchy is lower than the one attained by non-homogeneous latencies, and it is tight even for parallel-links topologies (Theorem 2), thus answering the open question posed by Roughgarden in [21].

To summarize, we obtain that the Price of Anarchy is independent of the network topology (i.e., the worst-case PoA is attained by parallel-link games) if and only if one of the following cases occurs: (i) $\theta = 0$ (which include the case of homogeneous latency functions as a special case) and (ii) $\theta = \infty$.

For the sake of a more concrete exposition of our results and for empirical purposes, we provide explicitly an instantiation of the PoA bounds in the case of polynomial latency functions. The resulting bounds depend on both the maximum and minimum degree of the polynomials and, in the case of non-homogeneous polynomials only, they also depend on θ . A quantitative representation of our results is partially summarized in Table 1.

Due to the lack of space, a detailed discussion of the results and some missing proofs are deferred to the full version of this paper [1].

Table 1. The Price of Anarchy of free-flow games with non-homogeneous (i.e., with constant terms allowed) polynomial latency functions of maximum degree $p \leq 4$ and minimum degree q . Unlabelled bounds are proven in this paper. Bounds for homogeneous (i.e., without constant terms) polynomials can be obtained from the case $\theta = 0$ (the same upper bounds have been given in [9], but tight lower bounds were only conjectured to exist). As it can be appreciated, the PoA depends on the network topology whenever $0 < \theta < \infty$.

(p, q)	$\theta = 0$		$\theta = 1/2$		$\theta = 1$		$\theta = \infty$	
	General	Path-disjoint	General	Path-disjoint	General	Path-disjoint	General	Path-disjoint
(1, 1)	1 [4]	1	1.1547 [4]	1.0909	1.2071 [4]	1.1429	1.3333 [23]	1.3333 [23]
(2, 1)	1.0355	1.0355	1.2873	1.1472	1.3852	1.2383	1.6258 [21]	1.6258 [21]
(2, 2)	1	1	1.2873	1.1472	1.3852	1.2383	1.6258 [21]	1.6258 [21]
(3, 1)	1.0982	1.0982	1.4078	1.1869	1.5475	1.3093	1.8956 [21]	1.8956 [21]
(3, 2)	1.0147	1.0147	1.4078	1.1869	1.5475	1.3093	1.8956 [21]	1.8956 [21]
(3, 3)	1	1	1.4078	1.1869	1.5475	1.3093	1.8956 [21]	1.8956 [21]
(4, 1)	1.1676	1.1676	1.5202	1.2170	1.6994	1.3652	2.1505 [21]	2.1505 [21]
(4, 2)	1.0450	1.0450	1.5202	1.2170	1.6994	1.3652	2.1505 [21]	2.1505 [21]
(4, 3)	1.0080	1.0080	1.5202	1.2170	1.6994	1.3652	2.1505 [21]	2.1505 [21]
(4, 4)	1	1	1.5202	1.2170	1.6994	1.3652	2.1505 [21]	2.1505 [21]

1.2 Related Work

Price of Anarchy in Routing Games: Introduced by Koutsoupias and Papadimitriou [13], the ratio between the social cost of the worst equilibrium of a game and its optimum was given the name Price of Anarchy (PoA) in [20]. For networks of linear latency and general topology, PoA was bounded tightly by $4/3$ [23] and $5/2$ in the atomic case [6]. Following results by Roughgarden [22] studied more general latency functions and atomic routing games and again gave tight bounds on PoA. However, for a large class of natural latency functions, PoA tends to 1 as the demand on the network approaches infinitesimally small or infinitely high levels [7, 8]. This casts doubts on the predictive power of PoA on the state of a real system, as noted in Monnot et al. [16].

Strategy sets of routing games are typically exponential in the number of vertices, hence restricting them is a common assumption. The unnatural character of Pigou in real systems was noted by Lu and Yu [15], who assume players have at least one strategy that is not more than λ away from the fastest strategy in congestion games. Restricting the strategy sets to obtain tighter bounds for PoA is also employed in [3], [?] for load balancing games (i.e., congestion games where the strategies of players are singleton sets). Fotakis [10] proved a pure PoA bound for symmetric atomic congestion games on extension-parallel networks, an interesting class of networks with linearly independent paths, that is equal to that of non-atomic congestion games.

Primal-dual techniques for bounding the Price of Anarchy in non-cooperative games have been proposed by Bilò [2], Kulkarni and Mirrokni [14], Nadav and Roughgarden [18] and Thang [24]. The methods proposed in [2] and [18] operate by explicitly formulating the problem of maximizing the Price of Anarchy of a class of games. Despite using the same formulation, they differ in the choice of

the variables. While [18] uses the probability distributions defining the outcomes occurring in the formulation, [2] adopts suitable multipliers for the resource cost functions. The methods in [14] and [24], instead, build on a formulation for the problem of optimizing the social function, and then implement the equilibria conditions within the choice of the dual variables. We adopt the method proposed in [2] as it appears to be more flexible and powerful in our realm of application. The first advantage is that it generalizes to any type of cost functions, while all the others require some restrictions: the method in [18] can only be applied to affine functions, the one in [14] requires convex functions, while that of [24] needs non-decreasing ones. Secondly, the method (if properly used) always yields tight bounds on the Price of Anarchy, while those in [14] and [24] are limited by the integrality gap of the formulation. Last but not least, it models in a simple, direct and intuitive way any new twist, as the free-flow property considered in this work, one may want to add to the scenario of application.

Transportation Research: The seminal work of Wardrop [26] introduces and formalizes one of the first notions of equilibrium in transportation networks. A proof of the equal social costs for equilibria and optimum (i.e., $\text{PoA} = 1$) in parallel links routing games appears in Nagurney and Qiang [19]. Related ideas from sensitivity analysis for edge cost functions are treated in Tobin and Friesz [25]. The Price of Anarchy was estimated for the city of Boston with different means from our study by Zhang et al. [27], where the sensitivity of the social cost at equilibrium with respect to edge parameters is also discussed. The previously cited works rely on the BPR estimation of cost functions [5], which are included in the family of weakly monomial latency functions we define in Sect. 3. The free-flow property in transportation networks has been first proposed by Jahn et al. [12] with respect to the problem of optimizing a centralized traffic flow without imposing too longer detours to some users.

2 Experimental Evidence for θ -Free-Flow Time in Singapore

We look for experimental evidence that commuters use the heuristic presented in the introduction to guide their routing decisions. Namely, we make the conjecture that commuters consider only paths with “length” at most a multiplicative factor $1 + \theta$ away from the shortest path taking them to their destination (where “length” is measured as a latency, or travel time). Does this conjecture hold in practice? To answer, we obtain data on the routing behavior of a sampled population. Modelling assumptions and a formal definition of θ are presented in the next Sect. 3.

2.1 The National Science Experiment

The NSE is a nationwide project in Singapore in which over 90,000 students from primary, secondary and junior college wore a sensor, called SENSg, for up to one

week per student in 2015 and 2016. The SENSg sensors collect various environmental data, and 9-degree of freedom motion data sampled every 13s using the Wi-Fi based localization system. The semantic data covers the identification of individual trips within the discrete stream of locations [11], inference of the activity performed at each endpoint and transportation mode classification [16, 17].

We use the NSE 2016 dataset which contains data from 49,526 students, and we implement the mode identification algorithm developed in [28] where five different modes can be identified, namely: (a) stationary; (b) walking; (c) riding a train; (d) riding a bus; and (e) riding a car. To ensure the quality of our empirical results, we perform a strict data cleaning process over the complete dataset. A total of 34,121 clean trips are considered, with 16,563 unique students and 89 different schools. This work focuses on morning travels of students who get to their schools from their homes.

Our dataset contains highly granular information concerning the routing decisions of the subjects. With the help of the onboard sensors in the device and the mode identification algorithm, we are able to obtain for each trip an accurate representation of its segments and their endpoints. For instance, typical segments making up a trip may be “Walk - Car - Walk”, or “Walk - Bus - Train - Bus - Walk”. The following study focuses on car trip segments. In this dataset [16], looking at the population of public transport users only, Price of Anarchy was upper bounded by 1.18. Conversely, Price of Anarchy for car users only was bounded by 1.86. Putting both populations together, Price of Anarchy was bounded by 1.34.

2.2 Estimation of Free-Flow Time for Selected Route

We compute a graph representation from a road map of Singapore, where each vertex is located at an intersection or a bend in the road. Every edge is assigned with a cost representing how much time is needed to traverse it. This latency is obtained from edge features such as the road type and the posted speed limit on the road. For each private transportation trip segment in the dataset, we associate its origin and destination with the closest vertex in the graph. We run a shortest path algorithm to estimate the free-flow travel time of the trip segment, referred to in the following as the *best free-flow time*. This best free-flow time is compared with the *data free-flow time*, or the time it would take the subject to travel its selected route if no one was on the road. We describe how the data free-flow time is estimated in the following paragraph.

A segment measured by the sensor consists of a stream of geographical locations. For each datapoint, we associate the closest edge in the graph. The size of the graph (61,151 vertices and 65,596 edges) implies a lengthy lookup phase to associate the point to its closest edge. For this reason, we consider a smaller dataset of 449 car segments out of the 17,897 segments in the larger dataset. These selected segments are well distributed across Singapore as depicted by Fig. 2. By adding the free-flow time of traversing each edge associated to the

data points, connected via heuristics detailed in our online version [1], we obtain the *data free-flow time*.

2.3 Estimate of θ

We compare the best free-flow time to the data free-flow time for each sample in our dataset, and denote by θ the percent increase between the two. A small value of θ yields support to the hypothesis that agents only consider routes which connect origin and destination in a straightforward manner (under no congestion) as part of their strategy set, see Fig. 3.




Fig. 3. The deviation is measured by the ratio of the selected route free-flow time to the minimum free-flow time among all routes between the origin and the destination. Close to 80% of the θ values are below 1, implying that the free-flow time of the selected route is rarely twice as long as the best free-flow time.

This experimental result provides justification for the upper bound of PoA estimated from the same dataset in previous work [16]. This benchmark is meaningful for real road networks, as latency functions are typically estimated using affine quartic monomials [5]. As noted in our introduction as well as in more details in the next section, our model is based on the assumption of a uniform θ bound over the whole population. We should note that this assumption is consistent with our experimental measurements, since these measurements provide us with estimates on the lower bounds of the agents’ θ ’s. More detailed models with a heterogeneous population/distribution of θ ’s is an interesting direction for future work.

3 Model and Definitions

For a positive integer i , let $[i] := \{1, 2, \dots, i\}$. Given a set A and a set $B \supseteq A$, let $\chi_A : B \rightarrow \{0, 1\}$ denote the indicator function, i.e., $\chi_A(x) = 1$ if $x \in A$ and $\chi_A(x) = 0$ if $x \notin A$. Given a tuple of numbers $(\alpha_1, \alpha_2, \dots, \alpha_k)$, we write $(\alpha_1, \alpha_2, \dots, \alpha_k) > 0$ if $\alpha_i \geq 0$ for any $i \in [k]$ and $\alpha_i > 0$ for some $i \in [k]$.

Non-atomic Congestion Games. A *non-atomic congestion game* (from now on, simply a *congestion game*) is a tuple $\text{CG} = ([n], (r_i)_{i \in [n]}, E, (\ell_e)_{e \in E}, (\Sigma_i)_{i \in [n]})$, where $[n]$ is a set of types, E is a set of resources, $\ell_e : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ is the *latency function* of resource $e \in E$, and, for each $i \in [n]$, $r_i \in \mathbb{R}_{\geq 0}$ is the

amount of players of type i and $\Sigma_i \subseteq 2^E \setminus \emptyset$ is the set of strategies for players of type i (i.e. a strategy is a non-empty subset of resources). We assume that latency functions are non-decreasing, positive, and continuous.

Classes of Congestion Games. A *network congestion game* is a congestion game based on a graph $G = (V, E)$, where the set of resources coincides with E , each type i is associated with a pair of nodes $(u_i, v_i) \in V \times V$, so that the set of strategies of players of type i is the set of paths from u_i to v_i in graph G . If there exists $u^* \in V$ such that $u^* = u_i$ for any $i \in [n]$, the game is called *single-source network congestion game*. Let \mathcal{P} be the set of all the paths P connecting source u_i with destination v_i , for any pair source-destination (u_i, v_i) . The game is called *path-disjoint network congestion game* if all the paths in \mathcal{P} are pair-wise node-disjoint.

A *load balancing game* is a congestion game in which each strategy is a singleton, i.e., $S = \{e\}$ for some $e \in E$, for any strategy $S \in \Sigma_i$ and type $i \in [n]$. A *parallel-link game* (or *symmetric load balancing game*) is a load balancing game in which all players have the same set of strategies. It is well-known that each load balancing game (resp. parallel-link game) can be modelled as a single-source congestion game (resp. path-disjoint network congestion game).

Latency Functions. For the sake of simplicity, we extend the domain of each latency function $\ell(x)$ to $x = 0$ in such a way that $\ell(0) = \lim_{x \rightarrow 0^+} \ell(x)$. Given a class of latency functions \mathcal{F} , let $[\mathcal{F}]_H := \{f : f(x) = g(x) - g(0), g \in \mathcal{F}\}$. Observe that $f(0) = 0$ for any $f \in [\mathcal{F}]_H$ by definition. In the following, we use similar definitions as in [21]. \mathcal{F} is *homogeneous* if $\mathcal{F} = [\mathcal{F}]_H$. \mathcal{F} is *weakly diverse* if $[\mathcal{F}]_H \subseteq \mathcal{F}$ and it contains at least one constant function (i.e., a function f such that $f(x) = \beta$ for any $x > 0$, for some $\beta > 0$). \mathcal{F} is *scale-closed* if it contains all the functions f such that $f(x) = \alpha g(x)$, for any $g \in \mathcal{F}$ and $\alpha > 0$. \mathcal{F} is *strongly diverse* if contains all the functions f such that $f(x) = \alpha g(x) + \beta$, for any $g \in [\mathcal{F}]_H$ and $(\alpha, \beta) > 0$. A *polynomial latency function* of maximum degree p and minimum degree q (with $p \geq q \geq 1$) is defined as $\ell_e(x) := \sum_{d=q}^p \alpha_{e,d} x^d + \beta_e$, where $\alpha_{e,q}, \alpha_{e,q+1}, \dots, \alpha_{e,p}, \beta_e > 0$. Let $\mathcal{P}_{p,q}$ denote the class of polynomial latency functions of maximum degree p and minimum (non-zero) degree q . A latency function ℓ_e is *affine* if $\ell_e \in \mathcal{PP}_1$.

Strategy Profiles and Pure Nash Equilibria. A *strategy profile* is a tuple $\sigma := (\sigma_{i,S})_{i \in [n], S \in \Sigma_i}$ with $\sum_{S \in \Sigma_i} \sigma_{i,S} = r_i$ for any $i \in [n]$, that is a state of the game where $\sigma_{i,S} \geq 0$ is the total amount of players of type i selecting strategy S for any $i \in [n]$ and $S \in \Sigma_i$. Given a strategy profile σ , $k_e(\sigma) := \sum_{i \in [n], S \in \Sigma_i : e \in S} \sigma_{i,S}$ is the *congestion* of e in σ , i.e., the total amount of players selecting e in σ , and given a strategy S , $c_S(\sigma) := \sum_{e \in S} \ell_e(k_e(\sigma))$ is the *cost* of players selecting S in σ . A strategy profile σ is a *pure Nash equilibrium* (or *Wardrop equilibrium*, or *equilibrium flow*) if and only if, for each $i \in [n]$, $S \in \Sigma_i : \sigma_{i,S} > 0$ and $S' \in \Sigma_i$, it holds that $c_S(\sigma) \leq c_{S'}(\sigma)$.

Quality of Equilibria. A *social function* that is usually used as a measure of the quality of a strategy profile in congestion games is the *total latency*, defined

as $\text{SUM}(\sigma) := \sum_{e \in E} k_e(\sigma) \ell_e(k_e(\sigma)) = \sum_{i \in [n]} r_i c_i(\sigma)$ at equilibrium σ . A *social optimum* is a strategy profile σ^* minimizing SUM.

The *Price of Anarchy* of a congestion game CG (with respect to the social function SUM), denoted as $\text{PoA}(\text{CG})$, is the supremum of the ratio $\text{SUM}(\sigma)/\text{SUM}(\sigma^*)$, where σ is a pure Nash equilibrium for CG and σ^* is a social optimum for CG. As shown in [23], all pure Nash equilibria of any congestion game have the same total latency. Thus, the Price of Anarchy can be redefined as the ratio $\text{SUM}(\sigma)/\text{SUM}(\sigma^*)$, where σ is an arbitrary pure Nash equilibrium for CG and σ^* is a social optimum for CG.

Free-Flow Congestion Games. Given $\theta \in [0, \infty]$, a θ -free-flow congestion game CG_θ is a congestion game in which, for each $i \in [n]$ and $S, S' \in \Sigma_i$, it holds that $\sum_{e \in S} \ell_e(0) \leq (1 + \theta) \sum_{e \in S'} \ell_e(0)$, i.e., all the strategies available to players of type i , when evaluated in absence of congestion, are within a factor $1 + \theta$ one from the other. Observe that free-flow congestion games are congestion games obeying some special properties. Thus, all positive results holding for congestion games carries over to θ -free-flow congestion games for any value of θ . Moreover, for $\theta = \infty$, any congestion game is a θ -free-flow congestion game.

4 Price of Anarchy of Free-Flow Congestion Games

In this section, we give tight bounds on the Price of Anarchy of free-flow congestion games. A detailed discussion of the implications of our theoretical results on the Price of Anarchy and how they relate to previous work, is given in the full version of this paper. Before going into details, we sketch the high level building blocks of the proofs of the upper bounds. For the general case, by adapting [2], we formulate the problem of bounding the Price of Anarchy of θ -free-flow congestion games by means of a factor-revealing pair of primal-dual linear programs. The techniques work as follows.

Given a θ -free-flow congestion game CG_θ and a family of latency functions \mathcal{F} , we know that we can model the latency of every resource $e \in E$ as $\ell_e(x) = \alpha_e f_e(x) + \beta_e$, with $f_e \in [\mathcal{F}]_H$, $\alpha_e \in \{0, 1\}$ and $\beta_e \geq 0$. We fix a Nash equilibrium σ and a social optimum σ^* for CG_θ . Hence, for every $e \in E$, the congestions $k_e(\sigma)$ and $k_e(\sigma^*)$ of e in σ and σ^* , respectively, become fixed constants. As the Price of Anarchy measures the worst-case ratio of $\text{SUM}(\sigma)$ over $\text{SUM}(\sigma^*)$, our goal becomes that of choosing suitable values for α_e and β_e , for every $e \in E$, so as to maximize $\text{SUM}(\sigma)$ under the assumption that $\text{SUM}(\sigma^*) = 1$, σ is a Nash equilibrium and CG_θ is a θ -free-flow game. In particular, constraint $\text{SUM}(\sigma^*) = 1$ can be assumed without loss of generality by a simple scaling argument, provided we relax the condition $\alpha_e \in \{0, 1\}$ with $\alpha_e \geq 0$. Thus, an optimal solution to the resulting linear program, call it LP, provides an upper bound to the Price of Anarchy of CG_θ . Next step is to compute and analyze the dual of LP, that we call DLP. DLP has three variables, namely x, y and γ , with $x \geq 0, y \geq 0$ and γ defining its objective value. Thus, by the Weak Duality Theorem, any feasible solution (x^*, y^*, γ^*) for DLP yields an upper bound of γ^* to the optimal solution of LP and so an upper bound to the Price of Anarchy of CG_θ . For each

function $f_e \in \mathcal{F}_H$, DLP has two constraints, namely $c_1(f_e, k_e(\sigma), k_e(\sigma^*), x, \gamma)$ and $c_2(f_e, k_e(\sigma), k_e(\sigma^*), y, \gamma)$, providing two lower bounds on γ , denoted as $\gamma(\mathcal{G}) := \inf_{x \geq 1} \sup_{l > 0, f \in \mathcal{G}} \left(\frac{k+x(-k+l)}{l} \right) \frac{f(k)}{f(l)}$ and $\gamma_\theta(\mathcal{G}) := \sup_{k > l > 0, f \in \mathcal{G}} \frac{f(k)(k(1+\theta)-l)}{f(k)(k-l)(1+\theta)+lf(l)\theta}$.

An important advantage of the primal-dual method is that, whenever LP provides a tight characterization of the properties possessed by the games and the equilibria under analysis, an optimal solution to DLP can be fruitfully exploited to construct, quite systematically, but not without effort, matching lower bounding instances. We manage to achieve this result also in this case, but, given the very technical nature of the constructions, we refer the interested reader to the full version of this paper.¹

For the case of parallel-links and path-disjoint games, we apply a similar, although more direct approach. We fix once again CG_θ , the family of latency functions \mathcal{F} , the latency of every resource $e \in E$, a Nash equilibrium σ and a social optimum σ^* for CG_θ , so as to obtain constant values for both $k_e(\sigma)$ and $k_e(\sigma^*)$. This time, instead of resorting to linear programming, we write down the parametric expression of the Price of Anarchy as a function of $k_e(\sigma)$, $k_e(\sigma^*)$ and the latency functions of the resources in the game. A key feature of this case, that makes it different from the general setting analyzed before, is that, here, we need have $\sum_{e \in E} k_e(\sigma) = \sum_{e \in E} k_e(\sigma^*)$. By exploiting this equality, together with the equilibrium conditions and the θ -free-flow property of CG_θ , we create a sequence of more and more relaxed upper bounds for the Price of Anarchy, until we end up to a sufficiently simple formula. Also in this case, we can show that the performed analysis is tight by providing matching lower bounding instances whose description is again deferred to the full version of this paper.

4.1 The Main Theorems

Theorem 1. *Let CG_θ be a θ -free-flow congestion game with latency functions in \mathcal{F} and $\theta \in [0, \infty]$. We have $\text{PoA}(\text{CG}_\theta) \leq \gamma([\mathcal{F}]_H)$ if $\theta = 0$, $\text{PoA}(\text{CG}_\theta) \leq \gamma(\mathcal{F})$ if $\theta = \infty$, and $\text{PoA}(\text{CG}_\theta) \leq \max\{\gamma([\mathcal{F}]_H), \gamma_\theta([\mathcal{F}]_H)\}$ if $\theta \in (0, \infty)$. These bounds are tight for single-source network games if \mathcal{F} is weakly diverse and even for load balancing games if \mathcal{F} is strongly diverse.*

We now show that, when considering either parallel-links games or path-disjoint network congestion games, a better bound on the Price of Anarchy can be achieved. To this aim, given a class of latency functions \mathcal{G} , let us define $\eta_\theta(\mathcal{G}) := \sup_{k > l > 0, f \in \mathcal{G}} \frac{kf(k)(1+\theta)}{kf(k)(1+\theta)+(lf(l)-lf(k))\theta}$.

Theorem 2. *Fix a value $\theta \in [0, \infty)$ and a class of latency functions \mathcal{F} . Let PLG_θ be a θ -free-flow path-disjoint network congestion game with latency functions in*

¹ In the related literature, bounds on the Price of Anarchy are often obtained by exploiting Roughgarden’s smoothness framework [22]. Similarities and differences between such framework and the primal-dual method are given in the full version of this paper [1].

\mathcal{F} . Then, $\text{PoA}(\text{PLG}_\theta) \leq \max\{\gamma([\mathcal{F}]_H), \eta_\theta([\mathcal{F}]_H)\}$. The bound is tight in general and even for parallel-links networks if \mathcal{F} is scale-closed.

By using Theorems 1 and 2, we can determine the exact Price of Anarchy of free-flow congestion games with polynomial latency functions in $\mathcal{P}_{p,q}$. In particular, we show that $\gamma_\theta([\mathcal{P}_{p,q}]_H) = \sup_{t>1} \frac{t^p(t(1+\theta)-1)}{t^p(t-1)(1+\theta)+\theta}$, $\gamma([\mathcal{P}_{p,q}]_H) = \frac{p^p}{(p+1)^{p+1}} \left(p^{-q} \sqrt{\left(\frac{(p+1)^{p+1} q^q}{(q+1)^{q+1} p^p} \right)^{p+1}} \right) \left(p^{-q} \sqrt{\frac{(p+1)^{p+1} q^q}{(q+1)^{q+1} p^p}} - 1 \right)^{-1} \chi_{[p-1]}(q) + \chi_{\{p\}}(q)$, $\eta_\theta([\mathcal{P}_{p,q}]_H) = \sup_{t>1} \frac{t^{p+1}(1+\theta)}{t^{p+1}(1+\theta)+(1-t^p)\theta}$, and by using such values in Theorems 1 and 2 we are able to derive tight bounds on the Price of Anarchy.

References

1. Benita, F., Bilò, V., Monnot, B., Piliouras, G., Vinci, C.: Data-driven models of selfish routing: why price of anarchy does depend on network topology. arXiv preprint [arXiv:2009.12871](https://arxiv.org/abs/2009.12871) (2020)
2. Bilò, V.: A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory Comput. Syst.* **62**(5), 1288–1317 (2018)
3. Bilò, V., Vinci, C.: On the impact of singleton strategies in congestion games. In: 25th Annual European Symposium on Algorithms (ESA 2017). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2017)
4. Bilò, V., Vinci, C.: The price of anarchy of affine congestion games with similar strategies. *Theor. Comput. Sci.* **806**, 641–654 (2020)
5. Bureau of Public Roads: Traffic assignment manual. US Department of Commerce (1964)
6. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, pp. 67–73. ACM (2005)
7. Colini-Baldeschi, R., Cominetti, R., Mertikopoulos, P., Scarsini, M.: The asymptotic behavior of the price of anarchy. In: Devanur, N.R., Lu, P. (eds.) WINE 2017. LNCS, vol. 10660, pp. 133–145. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71924-5_10
8. Colini-Baldeschi, R., Cominetti, R., Scarsini, M.: On the price of anarchy of highly congested nonatomic network games. In: Gairing, M., Savani, R. (eds.) SAGT 2016. LNCS, vol. 9928, pp. 117–128. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53354-3_10
9. Dumrauf, D., Gairing, M.: Price of anarchy for polynomial wardrop games. In: Spirakis, P., Mavronicolas, M., Kontogiannis, S. (eds.) WINE 2006. LNCS, vol. 4286, pp. 319–330. Springer, Heidelberg (2006). https://doi.org/10.1007/11944874_29
10. Fotakis, D.: Congestion games with linearly independent paths: convergence time and price of anarchy. *Theory Comput. Syst.* **47**(1), 113–136 (2010)
11. Gemici, K., Koutsoupias, E., Monnot, B., Papadimitriou, C.H., Piliouras, G.: Wealth inequality and the price of anarchy. In: 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019)
12. Jahn, O., Möhring, R.H., Schulz, A.S., Moses, N.E.S.: System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* **53**(4), 600–616 (2005)

13. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49116-3_38
14. Kulkarni, J., Mirrokni, V.S.: Robust price of anarchy bounds via LP and fenchel duality. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015), pp. 1030–1049. SIAM (2015)
15. Lu, P.Y., Yu, C.Y.: Worst-case Nash equilibria in restricted routing. *J. Comput. Sci. Technol.* **27**(4), 710–717 (2012)
16. Monnot, B., Benita, F., Piliouras, G.: Routing games in the wild: efficiency, equilibration and regret. In: Devanur, N.R., Lu, P. (eds.) WINE 2017. LNCS, vol. 10660, pp. 340–353. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71924-5_24
17. Monnot, B., et al.: Inferring activities and optimal trips: lessons from Singapore’s national science experiment. In: Cardin, M.-A., Fong, S.H., Krob, D., Lui, P.C., Tan, Y.H. (eds.) Complex Systems Design & Management Asia. AISC, vol. 426, pp. 247–264. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29643-2_19
18. Nadav, U., Roughgarden, T.: The limits of smoothness: a primal-dual framework for price of anarchy bounds. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 319–326. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17572-5_26
19. Nagurney, A., Qiang, Q.: A relative total cost index for the evaluation of transportation network robustness in the presence of degradable links and alternative travel behavior. *Int. Trans. Oper. Res.* **16**(1), 49–67 (2009)
20. Papadimitriou, C.: Algorithms, games, and the internet. In: Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing, pp. 749–753. ACM (2001)
21. Roughgarden, T.: The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.* **67**(2), 341–364 (2003)
22. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *J. ACM (JACM)* **62**(5), 32 (2015)
23. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM (JACM)* **49**(2), 236–259 (2002)
24. Thang, N.K.: Game efficiency through linear programming duality. In: Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS 2019), vol. LIPIcs 124, pp. 66:1–66:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
25. Tobin, R.L., Friesz, T.L.: Sensitivity analysis for equilibrium network flow. *Transp. Sci.* **22**(4), 242–250 (1988)
26. Wardrop, J.: Some Theoretical Aspects of Road Traffic Research. Road paper, Institution of Civil Engineers (1952). <https://books.google.it/books?id=9zEpAQAAAJ>
27. Zhang, J., Pourazarm, S., Cassandras, C.G., Paschalidis, I.C.: The price of anarchy in transportation networks: data-driven evaluation and reduction strategies. *Proc. IEEE* **106**(4), 538–553 (2018)
28. Zhou, Y., Wang, J., Shi, P., Dahlmeier, D., Tippenhauer, N., Wilhelm, E.: Power-saving transportation mode identification for large-scale applications. arXiv preprint [arXiv:1701.05768](https://arxiv.org/abs/1701.05768) (2017)