

Reconciling Selfish Routing with Social Good

Soumya Basu^(✉), Ger Yang, Thanasis Lianas, Evdokia Nikolova,
and Yitao Chen

The University of Texas at Austin, Austin, USA
basusoumya@utexas.edu

Abstract. Selfish routing is a central problem in algorithmic game theory, with one of the principal applications being that of routing in road networks. Inspired by the emergence of routing technologies and autonomous driving, we revisit selfish routing and consider three possible outcomes of it: (i) θ -Positive Nash Equilibrium flow, where every path that has non-zero flow on all of its edges has cost no greater than θ times the cost of any other path, (ii) θ -Used Nash Equilibrium flow, where every used path that appears in the path flow decomposition has cost no greater than θ times the cost of any other path, and (iii) θ -Envy Free flow, where every path that appears in the path flow decomposition has cost no greater than θ times the cost of any other path in the path flow decomposition. We first examine the relations of these outcomes among each other and then measure their possible impact on the network's performance. Right after, we examine the computational complexity of finding such flows of minimum social cost and give a range for θ for which this task is easy and a range of θ for which this task is NP-hard for the concepts of θ -Used Nash Equilibrium flow and θ -Envy Free flow. Finally, we propose strategies which, in a worst-case approach, can be used by a central planner in order to provide good θ -flows.

1 Introduction

Two Sides of the Coin: Social Welfare vs Selfishness. A fundamental problem arising in the management of road-traffic and communication networks is routing traffic to optimize network performance. In the setting of road-traffic networks the average delay incurred by a unit of flow quantifies the cost of a routing assignment. From a collective perspective minimizing the average cost translates to maximizing the welfare obtained by society. Starting from the seminal works of Wardrop [24] and Beckman et al. [2], the literature on network games has differentiated between (1) the objective of a central planner to minimize average cost and thus find a socially optimal (SO) flow, and (2) the selfish objectives of users minimizing their respective costs. In the latter case, the network users acting in their own interest are assumed to converge to a Nash Equilibrium (NE) flow as further rerouting fails to improve their own objective.

The tension between the central planner and individual users has been an object of intense study and solutions such as toll placement or Stackelberg routing (e.g., [2, 15, 19]) have been proposed in the past, each facing criticism in terms

of implementation and fairness towards various users. To mitigate this tension in a way that is more fair to the users, we set out to explore the properties of alternative solution concepts where users under some reasonable incentive condition adopt a more “socially desirable” routing of traffic in between the Nash equilibrium (which has high social cost) and the social optimum (which may be undesirable/unfair to users on the longer paths) [20]. The advent of routing applications and the growing dependence of users on these applications places us at an epoch when such new ideas in mechanism design may be more relevant and also more readily integrated to practice.

Consider the scenario where some routing application presents the uninformed users with routes alongside the guarantees of “relative fairness” and “reasonable delay” and the users adopt the paths. In this scenario, one may naturally bring forth the questions of whether there exist solutions (flows) where good social welfare is achieved under an appropriate incentive condition for the users and if such solutions can be efficiently computed. An example of such a solution could be enforcing a θ -approximate Nash equilibrium of low social cost, where users are guaranteed to get assigned a path of cost no greater than θ times the cost of the shortest path and as such, the solution is “relatively fair”.¹ Yet, other solution concepts seem to arise naturally and are introduced below.

Selfishness and Envy. To achieve the coveted middle ground between the social optimum and Nash equilibrium, by combining good social welfare with satisfied users, we consider equilibria notions related to: (1) selfishness and (2) envy. First, we consider selfishness where users tend to selfishly improve their own cost. Here we make the distinction between *positive* paths, i.e. paths that have positive flow in all of their edges (note, this is independent of the path flow decomposition), and *used* paths, i.e. paths that appear in the path flow decomposition with positive flow. With these definitions we define a θ -Positive Nash Equilibrium (θ -PNE)² to be a flow in which the length of any *positive* path in the network is less than or equal to θ times the length of any other path, and a θ -Used Nash Equilibrium (θ -UNE) to be a flow in which the length of any *used* path in the network is less than or equal to θ times the length of any other path. As we shall see, the set of θ -PNE flows is in general a strict subset of the set of θ -UNE flows, though for $\theta = 1$ these sets coincide. The definition of θ -approximate Nash equilibrium [9] corresponds to that of θ -UNE.

Next, we consider the notion of envy where for the same source and destination a user experiences envy against another user if the latter incurs smaller delay compared to the former under a given path flow. Similarly to the approximate Nash equilibrium flow we can consider a notion of approximately envy free flows where in a θ -Envy Free (θ -EF) flow, the ratio of any two *used* paths in

¹ The concept of fairness in selfish routing has been considered in the past, with the two main approaches defining fairness as: (1) the ratio of the maximum path delay in a given flow to the average delay under Nash equilibrium [20] and (2) the ratio of the maximum path delay to the minimum path delay in a given flow [11].

² In the literature, PNE is typically used for abbreviating Pure Nash Equilibrium, but we always use it to denote Positive Nash Equilibrium, as defined here.

the network is upper bounded by θ , for some $\theta \geq 1$. Note, the difference from the θ -UNE definition is that a used path's cost is compared only to other used paths' costs. Envy free flows arise naturally once we consider the routing applications setup where users only collect information about the routes provided by the application. On the one hand, the possible costs for the current users in some sense compare to the costs of the users that have already used the network. On the other hand, routes for which there is no (sufficient) information, i.e., routes that have not been chosen in the past (sufficiently many times), potentially may never appear as an option. Another motivation for a θ -EF flow arises from the literature on imitation games, e.g. [14], where users imitate other users with lower delay and jointly reach a fixed point which is a 1-EF flow.

An example of how the concepts of θ -PNE, θ -UNE, and θ -EF may differ from each other is illustrated in Fig. 1, with the details discussed in Sect. 3, where these notions are formally introduced.

Related Work. Starting from the seminal work of Koutsoupias and Papadimitriou [18], quantifying the worst case inefficiency of various non-cooperative games, including routing games, quickly became an intense area of research. In a routing game with arbitrary latency functions the ratio between the cost of a Nash equilibrium (NE) flow to the cost of a socially optimal (SO) flow may grow unbounded, as shown by Roughgarden and Tardos [22]. A series of papers have focused on developing techniques for bounding the inefficiency of the NE flow (e.g., [12, 16, 22]).

Considering the generalization to approximate NE flows, Caragiannis et al. [5–7] provided existential and computational results regarding approximate equilibria in weighted and unweighted atomic congestion games. Feldmann et al. [13] also considered computational issues for approximate equilibria and applied the method of randomized rounding to analyze which approximation guarantees can be achieved for atomic congestion games with latency functions in specific classes. Chen and Sinclair [8], focusing again on atomic congestion games, studied questions related to convergence times to approximate Nash equilibria. Christodoulou et al. [9] studied the performance of approximate Nash equilibria for atomic and non-atomic congestion games with polynomial latency functions by considering how much the price of anarchy worsens and how much the price of stability improves as a function of the approximation factor θ .

In a related thread of research, Jahn et al. [17] formalized the notion of constrained system optimal, where additional constraints were added along with the flow feasibility constraints. The additional constraints were introduced to reduce the unfairness of the resulting flow. Further, useful insights were obtained by Schulz and Stier-Moses [23] about the social welfare and fairness of these constrained system optimal flows. Recently, there have been efforts [3, 4] in quantifying the inefficiency needed to guarantee fairness among users. The authors there defined the ‘price of fairness’ as the proportional decrease of utility under fair resource allocation. As mentioned earlier, in routing games the fairness of socially optimal flows under different but related definitions has been studied by Roughgarden [20] and Correa et al. [10]. Further, Correa et al. [11] considered

the fairness and efficiency of min-max flows, where the objective is to minimize the maximum length of any used path in the network and noted how different path flows affect the fairness in the network even when the induced edge flows are identical.

Contribution. When users ask their routing devices for good origin to destination paths, they care about the end-to-end delay on their paths without (directly) caring about local (subpath) optimality conditions. This highlights that the path flows may play a key role in achieving the full potential for route planning mechanisms. On the conceptual side, through the study of the proposed solution concepts, i.e., the θ -PNE, θ -UNE and θ -EF, we clearly differentiate path flows from edge flows and study their varied effects in the balance between fairness and social cost.

On the technical side, we start by observing that the 1-UNE and the 1-PNE are indeed identical, which explains why the ‘used’ and the ‘positive’ paths have not been explicitly differentiated before this work. Beyond the case of $\theta = 1$, we notice that θ -PNE, θ -UNE and θ -EF flows are progressively larger sets, each containing the previous one, with promise of better tradeoff between the social welfare and fairness. In order to grasp the large separation between these concepts note that for some networks the θ -UNE is not contained in $\Omega(n\theta)$ -PNE, where n is the number of nodes in the network (Lemma 1).

Motivated from the classical study of the price of anarchy (PoA) of equilibrium flows we investigate the PoA of θ -UNE and θ -EF. In general we expect that as we move from θ -PNE to θ -EF flows, from a worst case perspective, we will encounter flows with larger social cost. As a worst case example we show that the PoA can be unbounded for 1-EF flows. However, the PoA upper bounds for both θ -PNE and θ -UNE turn out to be identical (Lemma 3). Our PoA bound generalizes the PoA bound of 1-PNE from [16]. Through a similar reasoning we show that the price of stability is non increasing from θ -PNE to θ -EF flows.

We next focus on computing a θ -PNE, a θ -UNE or a θ -EF flow with low social cost. The convex optimization approach for computing a socially optimal flow fails due to the non-convexity of the sets of θ -PNE, θ -UNE and θ -EF flows for $\theta > 1$ (Proposition 1). Formally, we prove (Theorem 1) that obtaining the best θ -UNE or the best θ -EF flow is NP-hard. Indeed given a socially optimal flow it is NP-hard to decide whether it admits a path flow decomposition which is θ -UNE (θ -EF) for arbitrarily large θ (assuming arbitrary latency functions). However, we leave open the complexity of finding the best θ -PNE flow ($\theta > 1$).

In the positive direction, we provide an approximation algorithm, based on a modified potential function, for designing a θ -PNE flow—which generates θ -EF and θ -UNE flows—with social cost guarantees. We explicitly derive the approximation ratio upper bound for solving minimization of social cost under the solution concepts for $\theta \geq 1$ for two classes of latency functions which are used in congestion networks, namely (1) polynomials with positive coefficients, (2) M/M/1 delay function (Theorem 3). This modified function approach was used by Christodoulou et al. [9] to derive upper bound for PoS(θ -UNE) with polynomial latency functions.

2 Preliminaries

Network and Flows. Consider a directed graph $G = (V, E)$ and a set of commodities \mathcal{K} with $K = |\mathcal{K}|$. Each commodity $k \in \mathcal{K}$ is associated with a source s_k , a sink t_k , and a demand $d_k > 0$. Each edge $e \in E$ is given a *latency function* $\ell_e(x)$, assumed to be standard, i.e. nonnegative, differentiable, and nondecreasing. We consider the standard *nonatomic network congestion game*, where each user routes an infinitesimal amount of flow. Let \mathcal{P}^k be the set of simple paths from s_k to t_k , and denote $\mathcal{P} = \cup_{k \in \mathcal{K}} \mathcal{P}^k$. A *feasible* flow can be represented as a *path flow vector* $\mathbf{f} = (f_\pi)_{\pi \in \mathcal{P}}$ that satisfies all demands, i.e. $\sum_{\pi \in \mathcal{P}^k} f_\pi = d_k$ for all $k \in \mathcal{K}$. The set of feasible path flow vectors is denoted by \mathcal{D}_p . A feasible path flow vector \mathbf{f} , induces a feasible *edge flow vector* in the network given as $\mathbf{x} = (x_e^k = \sum_{\pi \in \mathcal{P}^k: e \in \pi} f_\pi)_{e \in E, k \in \mathcal{K}}$. The congestion through edge e is the aggregate flow $x_e = \sum_{k \in \mathcal{K}} x_e^k$, for all $e \in E$. There may exist multiple feasible path flows, denoted as the set $\mathcal{D}_p(\mathbf{x})$, that give the same edge flow \mathbf{x} . We denote the set of feasible edge flows by \mathcal{D}_E .

Used and Positive Paths. Given a commodity $k \in \mathcal{K}$ and an edge flow vector \mathbf{x} , path $\pi \in \mathcal{P}^k$ is *positive* for commodity k if for all edges $e \in \pi$, $x_e^k > 0$. (Edge flow x_e is insufficient for this definition.) Given a commodity $k \in \mathcal{K}$ and a path flow vector \mathbf{f} , path $\pi \in \mathcal{P}^k$ is *used* by commodity k if $f_\pi > 0$ and *unused* otherwise. For each commodity $k \in \mathcal{K}$, we define the set of *positive* paths under edge flow \mathbf{x} as $\mathcal{P}_+^k(\mathbf{x})$ and the set of *used* paths under path flow \mathbf{f} as $\mathcal{P}_u^k(\mathbf{f})$. Note that a used path is always positive but a positive path may be unused.

Costs and Socially Optimal Flow. Under a path flow $\mathbf{f} \in \mathcal{D}_p$, the cost (latency) of a path π is defined to be the sum of latencies of edges along the path: $\ell_\pi(\mathbf{f}) = \ell_\pi(\mathbf{x}) = \sum_{e \in \pi} \ell_e(x_e)$ for $\mathbf{f} \in \mathcal{D}_p(\mathbf{x})$. The *social cost* (SC) of an edge flow $\mathbf{x} \in \mathcal{D}_E$ is $SC(\mathbf{x}) = \sum_{e \in E} x_e \ell_e(x_e)$. The social cost of a path flow is the social cost of its corresponding edge flow. The socially optimal edge/path flow is the flow that minimizes social cost among all feasible edge/path flows. The set of socially optimal edge/path flows is denoted by $SO_E = \{\mathbf{x} \in \arg \min SC(\mathbf{x})\}$ and $SO_p = \{\mathbf{f} \in \arg \min SC(\mathbf{f})\}$, respectively.

Nash Equilibrium and Efficiency. A path flow \mathbf{f} is a *Nash Equilibrium* if for any commodity $k \in \mathcal{K}$ and any used path $p \in \mathcal{P}_u^k(\mathbf{f})$ we have $\ell_p(\mathbf{f}) \leq \ell_q(\mathbf{f})$, for all paths $q \in \mathcal{P}^k$. The efficiency of an equilibrium is often measured via the *price of anarchy* and the *price of stability*. Here we generalize them for arbitrary set of flows as in the following definitions. Given a set of path flows \mathcal{F} and socially optimal edge flow $\mathbf{x}^* \in SO_E$, the price of anarchy (PoA) and the price of stability (PoS) are defined as

$$PoA(\mathcal{F}) = \max \left\{ \frac{SC(\mathbf{f})}{SC(\mathbf{x}^*)} : \mathbf{f} \in \mathcal{F} \right\}, \quad PoS(\mathcal{F}) = \min \left\{ \frac{SC(\mathbf{f})}{SC(\mathbf{x}^*)} : \mathbf{f} \in \mathcal{F} \right\}.$$

3 Solution Concepts

Here we give the formal definition of the solution concepts we introduced in Sect. 1. We also provide an example to illustrate their differences, and prove

that each solution concept may correspond to a non-convex set of flows. We begin with the definitions of the solution concepts:

Definition 1 (θ -PNE, θ -UNE, and θ -EF).

1. An edge flow \mathbf{x} is a θ -Positive Nash Equilibrium (θ -PNE) flow if for any commodity $k \in \mathcal{K}$, any positive path $p \in \mathcal{P}_+^k(\mathbf{x})$, and all paths $q \in \mathcal{P}^k$, we have $\ell_p(\mathbf{x}) \leq \theta \ell_q(\mathbf{x})$.
2. A path flow \mathbf{f} is a θ -Used Nash Equilibrium (θ -UNE) flow if for any commodity $k \in \mathcal{K}$, any used path $p \in \mathcal{P}_u^k(\mathbf{f})$, and all paths $q \in \mathcal{P}^k$, we have $\ell_p(\mathbf{f}) \leq \theta \ell_q(\mathbf{f})$.
3. A path flow \mathbf{f} is θ -Envy Free (θ -EF) if for any commodity $k \in \mathcal{K}$, any used path $p \in \mathcal{P}_u^k(\mathbf{f})$, and all used paths $q \in \mathcal{P}_u^k(\mathbf{f})$, we have $\ell_p(\mathbf{f}) \leq \theta \ell_q(\mathbf{f})$.

For simplicity, we use θ -PNE, θ -UNE or θ -EF to describe the set of θ -PNE, θ -UNE or θ -EF flows, respectively. Also, we refer to them as θ -fair flows. To see how these concepts may differ from each other, we give an example in Fig. 1.

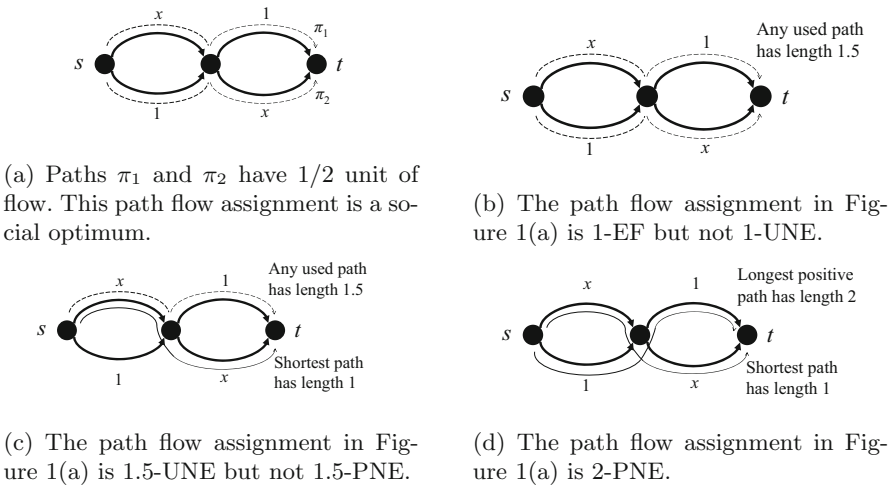



Fig. 1. Example illustrating the three solution concepts θ -UNE, θ -PNE and θ -EF.

Our goal is to examine the properties of θ -fair flows and provide ways to obtain such flows with good social cost. Regarding the second direction, in general, the sets of θ -PNE, θ -UNE, and θ -EF flows may not be convex and may contain multiple path flows, which raises the level of difficulty for computing good or optimal such flows. Next, we present an example that demonstrates the non-convexity of these sets (Fig. 2).

Proposition 1 (Non-convexity of θ flows). *There exists an instance such that the sets θ -PNE, θ -UNE, and θ -EF are not convex, for some $\theta > 1$.*



f_1 : s - u - v - t with flow 1	f_3 : s - u - v - t with flow 1/2
f_2 : s - u - t with flow 2/3	s - u - t with flow 1/3
s - v - t with flow 1/3	s - v - t with flow 1/6

Fig. 2. Non-convexity of θ -flows. Both f_1 and f_2 are $3/2$ -PNE/UNE/EF, but their convex combination (with even weights) f_3 is not.

The following two lemmas establish the hierarchy among the proposed solution concepts by showing a crisp containment of various flows. Due to space constraints, the proofs are presented in the full version [1], Sect. 4.

Lemma 1 (Hierarchy of θ -flows). *Given a multi-commodity network, for $\theta' > \theta \geq 1$, $\mathcal{F} \in \{\text{PNE}, \text{UNE}, \text{EF}\}$ satisfies $\theta\text{-}\mathcal{F} \subseteq \theta'\text{-}\mathcal{F}$. Further, for any $\theta \geq 1$, $\theta\text{-PNE} \subseteq \theta\text{-UNE} \subseteq \theta\text{-EF}$ holds. On the other hand, for any $\theta \geq 1$, there exists a network such that $1\text{-EF} \not\subseteq \theta\text{-UNE}$.*

In the following lemma, we further demonstrate the relationship between the θ -UNE and θ -PNE. We can see that 1-UNE and 1-PNE both coincide with the familiar Nash equilibrium.

Lemma 2 (θ -UNE and θ -PNE). *Given a multi-commodity network with n nodes, for any path flow f and its induced edge flow x , $f \in 1\text{-UNE}$ if and only if $x \in 1\text{-PNE}$. Further, for any $\theta > 1$, $\theta\text{-UNE} \subset ((n-1)\theta)\text{-PNE}$ holds. On the other hand, for any $\theta \geq 1.5$, there exists a network such that $\theta\text{-UNE} \not\subseteq ((n-3)\theta/3)\text{-PNE}$.*

We next analyze the cost of the θ -flows. Note that the θ flows are not unique for $\theta > 1$ and this implies that potentially under each solution concept we can have a range of attainable costs. From the containment relations of the θ flows (Lemma 1, Part 2), it follows that for any $\theta \geq 1$,

$$\begin{aligned} \text{PoA}(\theta\text{-EF}) &\geq \text{PoA}(\theta\text{-UNE}) \geq \text{PoA}(\theta\text{-PNE}), \\ \text{PoS}(\theta\text{-EF}) &\leq \text{PoS}(\theta\text{-UNE}) \leq \text{PoS}(\theta\text{-PNE}). \end{aligned}$$

Further, we present upper bounds on the PoA for θ -UNE and θ -PNE flows, and show that the PoA for 1-EF flow is unbounded. In that effort, we generalize techniques presented in [16] which was used for bounding $\text{PoA}(1\text{-PNE})$. We need the following definitions in order to bound PoA:

$$\omega(\mathcal{L}, \lambda) = \sup_{\ell \in \mathcal{L}} \sup_{x, x' \geq 0} \frac{(\ell(x) - \lambda \ell(x')) x'}{x \ell(x)}, \quad A(\theta) = \{\lambda \in \mathbb{R}^+ : \omega(\mathcal{L}, \lambda) \leq 1/\theta\}.$$

The following lemma summarizes the PoA results for the solution hierarchies. For the proof refer to the full version [1], Sect. 5.

Lemma 3 (PoA of θ -flows). *For latency functions in class \mathcal{L} , $PoA(\theta\text{-UNE}) \leq \inf_{\lambda \in \Lambda(\theta)} \theta \lambda (1 - \theta \omega(\mathcal{L}, \lambda))^{-1}$. On the other hand, there exists a network with linear latency functions for which the $PoA(1\text{-EF})$ is unbounded.*

4 Optimal θ -Flows: Complexity and Approximation

In this section, we first discuss the possibility of designing a flow that balances the fairness and the social cost in the network under the new solution concepts. The standard convex optimization approaches fail to find socially optimal θ flow as the sets of θ -PNE, θ -UNE and θ -EF flows are all non-convex. We formally prove that finding socially optimal θ -UNE or θ -EF flows is NP hard. Then, using a modified potential function technique we provide approximation guarantees for two common classes of latency functions used in congestion network modeling, namely (1) polynomial and (2) M/M/1.

Consider the instance in Fig. 3. The next proposition states that though the socially optimal flow—uniquely determined by the edge flow—is unfair in the worst case, there exist path flows which are fair or almost fair under the concepts of UNE and EF flows. The proof is in the full version [1], Sect. 7.1.




Fig. 3. Improved balance: example.

Proposition 2 (Balanced path flows). *For the n -stage instance depicted in Fig. 3 with $n\epsilon = 2$, the socially optimal edge flow is a 2-PNE. Moreover, the socially optimal flow admits path flows which are 1-EF or $(1 + 1/n)$ -UNE.*

4.1 Existence and Complexity

The previous motivating example naturally leads to the following computational problems given a $\theta \geq 1$.

- (P1) Find a θ -EF path flow with the minimal social cost.
- (P2) Find a θ -UNE path flow with the minimal social cost.
- (P3) Find a θ -PNE edge flow with the minimal social cost.

Existence of Polynomial-Size Solutions. An observation to Problem (P1) and (P2) is that the outputs of these two problems are path flow vectors, which are potentially of exponential size relative to the problem instances. The following lemma proves the existence of polynomial sized path flows, in absence of which there is no hope to find a polynomial time algorithm for these problems.

Lemma 4 (Existence of polynomial-size solutions). *Given a θ -EF (or a θ -UNE) path flow \mathbf{f} , there exists a θ -EF (resp., a θ -UNE) path flow \mathbf{f}' that uses at most $|E|$ paths for each source-sink pair and has the same edge flow as \mathbf{f} .*

Computational Complexity. We show that for large θ , the socially optimal flow is guaranteed to be contained in those θ -flows, and hence the optimal θ -flows can be computed efficiently. However, for small θ , we will show that solving Problem (P1) and Problem (P2) is NP-hard, while it remains open whether Problem (P3) can be computed efficiently. More precisely, for a latency class \mathcal{L} , this particular threshold is $\gamma(\mathcal{L}) = \min\{\gamma : \ell^*(x) \leq \gamma \ell(x), \forall \ell \in \mathcal{L}, \forall x \geq 0\}$, where $\ell^*(x) = \ell(x) + x\ell'(x)$. The main result of this section is:

Theorem 1 (Computational Complexity of (P1)–(P3)). *For any multi commodity instance with latency functions in any class \mathcal{L} , there is a polynomial time algorithm for solving Problem (P1)–(P3) for $\theta \geq \gamma(\mathcal{L})$. On the other hand, it is NP-hard to solve Problem (P1) for $\theta \in [1, \gamma(\mathcal{L})]$ and Problem (P2) for $\theta \in (1, \gamma(\mathcal{L}))$, for arbitrary single commodity instances with latency functions in an arbitrary class \mathcal{L} .*

The first part of Theorem 1 follows easily from the following lemma in [11].

Lemma 5 (P1)–(P3) for $\theta \geq \gamma(\mathcal{L})$ are easy [11]. *For an instance with latency functions in \mathcal{L} , any socially optimal path flow $\mathbf{f} \in SO_p$ is $\gamma(\mathcal{L})$ -PNE.*

For the proof of the second part of Theorem 1, we consider the class of polynomial functions of degree at most p , denoted by \mathcal{L}_p . We note that $\gamma(\mathcal{L}_p) = p+1$. We show that when the latency functions are in \mathcal{L}_p , then the related decision problems we state in Theorem 2, stated below, have polynomial-time reductions from the NP-complete problem PARTITION.

Theorem 2 (NP-hardness of (P1) and (P2)). *For an arbitrary single commodity instance with latency functions in class \mathcal{L}_p for $p \geq 1$, it is NP-hard to*

1. *decide whether a socially optimal flow has a θ -UNE path flow decomposition for $\theta \in (1, p+1)$.*
2. *decide whether a socially optimal flow has a θ' -EF path flow decomposition for $\theta' \in [1, p+1)$.*

Proof. (Proof Sketch) For lack of space we present a proof sketch mentioning the flow of key ideas behind the proof. The proof is divided into two parts. For the first part, we show the NP-hardness for 1.5-UNE and 1-EF path flow decompositions under the social optimum in Lemma 6, based on the construction in Theorem 3.3 in Correa et al. [11]. Then, in the second part, we propose a novel way to generalize the construction to the entire range of θ and θ' specified in Theorem 2.

Construction. Let $G(q)$ be the two-link parallel network with the top link e_u having latency $\ell_u(x) = q$ and the bottom link e_b having latency $\ell_b(x) = qx$. Given an instance of the PARTITION problem, $q_1, \dots, q_n, \sum_{i=1}^n q_i = 2B$, we now construct a single commodity network as the two-link n -stage network G , shown in Fig. 4. In stage i we connect $G(q_{i-1})$ to $G(q_i)$ to the right for $i = 2$ to n . A unit demand has to be routed from the source in $G(q_1)$ to the destination in $G(q_n)$. Finally, we augment to the right of G a two-link parallel network G' . For G' the top link latency is $\ell_{u,(n+1)}(x) = ax^p + b$ and bottom link latency is $\ell_{d,(n+1)}(x) = cx^p$. We set $a = \frac{\alpha B}{(1-3/8B)^p}$, $b = \beta B(p + 1)$, and $c = \frac{(\alpha + \beta)B}{(3/8B)^p}$, where $\alpha, \beta > 0$ are appropriate parameters (specified in the proof of Theorem 1 in the full version [1]). We call the entire network H .




Fig. 4. An instance of congestion game constructed from a given instance of PARTITION.

Lemma 6 states that it is NP-hard to find 1.5-UNE and 1-EF path flow decompositions under the social optimum.

Lemma 6 (Hardness Result in G). *For single commodity instances with linear latency functions it is NP-hard to decide whether a social optimum flow has a 1.5-UNE flow decomposition or a 1-EF flow decomposition.*

Amplification of Hardness. The next step is to amplify the hardness result to all $\theta \in (1, p + 1)$ for UNE and to all $\theta' \in [1, p + 1)$ for EF flows. The key observation that facilitates this amplification is the following claim.

Claim 1. If the answer to PARTITION is NO then in the sub-network G any path decomposition of the socially optimal flow o routes at least $\frac{1}{2B}$ amount of flow through paths of length strictly greater than $\frac{3}{2}B$.

Through careful combination of the paths it is shown that the socially optimal flow is a c_1 -UNE and c_2 -EF flow if and only if the given PARTITION instance has a YES answer. Here c_1 and c_2 are constants given by

$$c_1 = \frac{\alpha + \beta + \beta p + \frac{3}{2}}{1 + \alpha + \beta} = 1 + \frac{\frac{1}{2} + \beta p}{1 + \alpha + \beta}, \quad c_2 = \frac{\alpha + \beta + \beta p + \frac{3}{2}}{2 + \alpha + \beta} = 1 + \frac{-\frac{1}{2} + \beta p}{2 + \alpha + \beta}.$$

4.2 Approximation Using Modified Potential Functions

In this section, we provide a modified potential function based approximation algorithm to problems (P1), (P2) and (P3). The idea of modified potential functions was introduced for bounding the PoS of approximate Nash equilibria in [9]. For a given θ this approach produces a θ -PNE and, due to containment, any feasible path flow corresponding to the edge flow will be in θ -UNE and θ -EF.

Algorithm 1. Modified Potential Algorithm

Input: Multi-commodity network \mathcal{G} , θ .

Output: Edge flow $\mathbf{x}_A \in \theta$ -PNE and path flow $\mathbf{f}_A \in \theta$ -UNE \cap θ -EF .

1: For all $e \in E$ choose $\phi_e(x) \in [\ell_e(s)/\theta, \ell_e(x)]$ (Specified later.)

2: Compute $\mathbf{x}_A = \operatorname{argmin}_{\mathbf{x} \in \mathcal{F}} \sum_{e \in E} \int_{x=0}^{x_e} \phi_e(x) dx$.

3: Compute any path decomposition \mathbf{f}_A of \mathbf{x}_A .

Theorem 3 characterizes the performance of Algorithm 1 for two important classes of latency functions which are used for modeling congestion networks—(1) Polynomial latency with positive coefficients, and (2) M/M/1 latency.

Theorem 3 (Performance of Algorithm 1). *Given a multi-commodity network \mathcal{G} and $\theta \geq 1$ the algorithm produces an edge flow that is θ -PNE and a path flow that is both θ -UNE and θ -EF.*

(1) **Polynomial Latency:** *Additionally, let the latency function $\ell_e(x) = \sum_{k=0}^p a_{e,k} x^k$, $a_{e,k} \geq 0$, for all $e \in E$ and some finite p . Algorithm 1 with $\phi_e(x) = \sum_{k=0}^p \zeta_k a_{e,k} x^k$, $\zeta_k = (1 + \min\{k, \theta - 1\})/\theta$ for all $e \in E, k \leq p$, is a $\left(\theta \left(1 - \frac{p}{1+p} \left(\frac{\theta}{1+p} \right)^{1/p} \right) \right)^{-1}$ -approximation algorithm for the problems (P1), (P2) and (P3).*

(2) **M/M/1 Latency:** *Additionally, let the latency function $\ell_e(x) = 1/(u_e - x)$ with $u_e > 0, \rho_e = d_{tot}/u_e$ and $\rho_{\max} = \max_{e \in E} \rho_e < 1$, for all $e \in E$. Algorithm 1 with $\phi_e(x) = 1/(a_e u_e - x)$, $a_e = \{0, 1 - \theta(1 - \rho)\}/\rho$ for all $e \in E$, is a $\frac{1}{2} \left(1 + \frac{1}{\sqrt{1 - \rho_{\max}(\theta)}} \right)$ -approximation algorithm for the problems (P1), (P2) and (P3), where $\rho_{\max}(\theta) = \max\{0, 1 - \theta(1 - \rho_{\max})\}$.*

The first part of Theorem 3 (i.e., the output being θ flows) follows from the idea, that a 1-PNE flow under the modified potential is a θ -PNE flow under the original latency functions. To prove the second part, we bound the inefficiency of the flow \mathbf{x}_A as the PoA(1-PNE) under the modified potential functions, which in turn serves as an approximation ratio for the minimization of social cost under θ -PNE (P3), θ -UNE (P2) and θ -EF (P1). Using proper functions $\phi_e(\cdot)$ along with the λ - μ smoothness framework [21] we strictly improve the approximation ratio from PoA(1-PNE). Recall the trivial solution—1-PNE flow which can be computed efficiently gives an approximation ratio of PoA(1-PNE). The choice of $\phi_e(\cdot)$ and the subsequent bounds for polynomial latencies were presented in [9],

but for upper bounding $\text{PoS}(\theta\text{-UNE})$. The detailed proofs are presented in the full version [1], Sect. 7.2.

Acknowledgements. This work was supported in part by NSF grant numbers CCF-1216103, CCF-1350823 and CCF-1331863. Part of the research was performed while a subset of the authors were at the Simons Institute in Berkeley, CA in Fall 2015.

References

1. Basu, S., Yang, G., Lianas, T., Nikolova, E., Chen, Y.: Reconciling selfish routing with social good. [arXiv:1707.00208](https://arxiv.org/abs/1707.00208) (2017)
2. Beckmann, M., McGuire, C., Winsten, C.B.: Studies in the economics of transportation. Technical report (1956)
3. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. *Oper. Res.* **59**(1), 17–31 (2011)
4. Bertsimas, D., Farias, V.F., Trichakis, N.: On the efficiency-fairness trade-off. *Manag. Sci.* **58**(12), 2234–2250 (2012)
5. Caragiannis, I., Fanelli, A., Gravin, N., Skopalik, A.: Computing approximate pure nash equilibria in congestion games. *ACM SIGecom Exch.* **11**(1), 26–29 (2012)
6. Caragiannis, I., Fanelli, A., Gravin, N., Skopalik, A.: Approximate pure nash equilibria in weighted congestion games: existence, efficient computation, and structure. *ACM Trans. Econ. Comput.* **3**(1), 2 (2015)
7. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006*. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006). doi:[10.1007/11786986_28](https://doi.org/10.1007/11786986_28)
8. Chien, S., Sinclair, A.: Convergence to approximate nash equilibria in congestion games. In: *SODA* (2007)
9. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the performance of approximate equilibria in congestion games. *Algorithmica* **61**(1), 116–140 (2011)
10. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Computational complexity, fairness, and the price of anarchy of the maximum latency problem. In: Bienstock, D., Nemhauser, G. (eds.) *IPCO 2004*. LNCS, vol. 3064, pp. 59–73. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-25960-2_5](https://doi.org/10.1007/978-3-540-25960-2_5)
11. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Fast, fair, and efficient flows in networks. *Oper. Res.* **55**(2), 215–225 (2007)
12. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: A geometric approach to the price of anarchy in nonatomic congestion games. *Games Econ. Behav.* **64**(2), 457–469 (2008)
13. Feldmann, A.E., Röglin, H., Vöcking, B.: Computing approximate nash equilibria in network congestion games. In: Shvartsman, A.A., Felber, P. (eds.) *SIROCCO 2008*. LNCS, vol. 5058, pp. 209–220. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-69355-0_18](https://doi.org/10.1007/978-3-540-69355-0_18)
14. Fischer, S., Vöcking, B.: On the evolution of selfish routing. In: Albers, S., Radzik, T. (eds.) *ESA 2004*. LNCS, vol. 3221, pp. 323–334. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-30140-0_30](https://doi.org/10.1007/978-3-540-30140-0_30)
15. Fleischer, L., Jain, K., Mahdian, M.: Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In: *FOCS* (2004)

16. Harks, T.: On the price of anarchy of network games with nonatomic and atomic players. Technical report, available at Optimization Online (2007)
17. Jahn, O., Möhring, R.H., Schulz, A.S., Stier-Moses, N.E.: System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* **53**(4), 600–616 (2005)
18. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999*. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). doi:[10.1007/3-540-49116-3_38](https://doi.org/10.1007/3-540-49116-3_38)
19. Roughgarden, T.: Stackelberg scheduling strategies. In: *STOC* (2001)
20. Roughgarden, T.: How unfair is optimal routing? In: *SODA* (2002)
21. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *J. ACM (JACM)* **62**(5), 32 (2015)
22. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM (JACM)* **49**(2), 236–259 (2002)
23. Schulz, A.S., Stier-Moses, N.E.: Efficiency and fairness of system-optimal routing with user constraints. *Networks* **48**(4), 223–234 (2006)
24. Wardrop, J.G.: *Some Theoretical Aspects of Road Traffic Research* (1952)