

Balancing Fairness and Efficiency in Traffic Routing via Interpolated Traffic Assignment

Devansh Jalota
Institute for Computational and
Mathematical Engineering, Stanford
University
Stanford, California, USA
djalota@stanford.edu

Kiril Solovey
Technion - Israel Institute of
Technology
Haifa, Israel
kirilsol@stanford.edu

Matthew Tsao
Department of Electrical Engineering,
Stanford University
Stanford, California, USA
mwtsao@stanford.edu

Stephen Zoepf
Lacuna AI
Palo Alto, California, USA
stephen.zoepf@lacuna.ai

Marco Pavone
Department of Aeronautics and
Astronautics, Stanford University
Stanford, California, USA
pavone@stanford.edu

ABSTRACT

System optimum (SO) routing, wherein the total travel time of all users is minimized, is a holy grail for transportation authorities. However, SO routing may discriminate against users who incur much larger travel times than others to achieve high system efficiency, i.e., low total travel times. To address the inherent unfairness of SO routing, we study the β -fair SO problem whose goal is to minimize the total travel time while guaranteeing a $\beta \geq 1$ level of unfairness, which specifies the maximum possible ratio between the travel times of different users with shared origins and destinations.

To obtain feasible solutions to the β -fair SO problem while achieving high system efficiency, we develop a new convex program, the Interpolated Traffic Assignment Problem (I-TAP), which interpolates between a fairness-promoting and an efficiency-promoting traffic-assignment objective. We evaluate the efficacy of I-TAP through theoretical bounds on the total system travel time and level of unfairness in terms of its interpolation parameter, as well as present a numerical comparison between I-TAP and a state-of-the-art algorithm on a range of transportation networks. The numerical results indicate that our approach is faster by several orders of magnitude as compared to the benchmark algorithm, while achieving higher system efficiency for all desirable levels of unfairness. We further leverage the structure of I-TAP to develop two pricing mechanisms to collectively enforce the I-TAP solution in the presence of selfish homogeneous and heterogeneous users, respectively, that independently choose routes to minimize their own travel costs. We mention that this is the first study of pricing in the context of fair routing for general road networks.

KEYWORDS

Traffic Assignment; Congestion Games; Nash Equilibria

ACM Reference Format:

Devansh Jalota, Kiril Solovey, Matthew Tsao, Stephen Zoepf, and Marco Pavone. 2022. Balancing Fairness and Efficiency in Traffic Routing via Interpolated Traffic Assignment. In *Proc. of the 21st International Conference*

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

on *Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 9 pages.

1 INTRODUCTION

Traffic congestion has soared in major urban centres across the world, leading to widespread environmental pollution and huge economic losses. In the US alone, almost 90 billion US dollars of losses are incurred every year, with commuters losing hundreds of hours due to traffic congestion [13]. A contributing factor to increasing road traffic is the often sub-optimal route selection by users due to the lack of centralized control [23, 24]. In particular, *selfish routing*, wherein users choose routes to minimize their travel times, results in a user equilibrium (UE) traffic pattern that is often far from the system optimum (SO) [27, 32]. To cope with the efficiency loss due to the selfishness of users, several methods including the control of a fraction of compliant users [26] and marginal cost tolls, where users pay for the externalities they impose on others, have been used to enforce the SO solution as a UE [20, 30].

While determining SO tolls is of fundamental theoretical importance, it is of limited practical interest [31] since SO traffic patterns are often unfair with some users incurring much larger travel times than others. This discrepancy among user travel times is referred to as *unfairness*, which, more formally, is the maximum possible ratio across all origin-destination (O-D) pairs of the travel time of a given user to the travel time of the fastest user between the same O-D pair. The unfairness of the SO solution can be quite high in real-world transportation networks, since users may spend nearly twice as much time as others travelling between the same O-D pair [16]. Moreover, a theoretical analysis established that the SO solution can even have unbounded unfairness [22].

The lack of consideration of user-specific travel times in the SO problem has led to the design of methods that aim to achieve a balance between the total travel time of a traffic assignment and the level of fairness that it provides. In a seminal work, Jahn et al. [16] introduced the Constrained System Optimum (CSO) to reduce the unfairness of traffic flows by bounding the ratio of the normal length of a path of a given user to the normal length of the shortest path for the same O-D pair. Here, *normal length* is any metric for an edge that is fixed *a priori* and is independent of the traffic flow, e.g.,

edge length or free-flow travel time. While many approaches to solve the CSO problem have been developed [1–3], they suffer from the limitation that the level of experienced unfairness in terms of user travel times can be much higher than the bound on the ratio of normal lengths that the CSO is guaranteed to satisfy. In addition to this drawback, the algorithmic approaches to solve the CSO problem are often computationally prohibitive and do not provide theoretical guarantees in terms of the resulting solution fairness and efficiency. Furthermore, it is unclear how to develop a pricing scheme to enforce such proposed traffic assignments in practice.

In this work, we study a problem analogous to CSO that differs in the problem’s unfairness constraints. In particular, we explicitly consider the unfairness in terms of user travel times as in [4], which, arguably, is a more accurate representation of user constraints as it accounts for costs that vary according to a traffic assignment. Our work further addresses the algorithmic concerns of existing approaches to solve fairness-constrained traffic routing problems by developing (i) a computationally-efficient approach that trades off efficiency and fairness in traffic routing, (ii) theoretical bounds to quantify the performance of our algorithm, and (iii) a pricing mechanism to enforce the resulting traffic assignment.

Contributions. We study the β -fair System Optimum (β -SO) problem, which involves minimizing the total travel time of users subject to unfairness constraints, where a $\beta \geq 1$ bound on unfairness specifies the maximum allowable ratio between the travel times of different users with shared origins and destinations.

We develop a simple yet effective approach for β -SO that involves solving a new convex program, the interpolated traffic assignment problem (I-TAP). I-TAP interpolates between the fair UE and efficient SO objectives to achieve a solution that is simultaneously fair and efficient. This allows us to approximate the β -SO problem as an unconstrained TAP, which can be solved quickly. We further present theoretical bounds on the total system travel time and unfairness level in terms of the interpolation parameter of I-TAP.

We then exploit the structure of I-TAP to develop two pricing schemes which enforce users to selfishly select the flows satisfying the β bound on unfairness computed through our approach. For homogeneous users with the same value of time we develop a natural marginal-cost pricing scheme. For heterogeneous users, we exploit a linear-programming method [12]. We mention that our work is the first to study road pricing in connection with fair routing for general road networks as opposed to, e.g., parallel networks.

Finally, we evaluate the performance of our approach on real-world transportation networks. The numerical results indicate significant computational savings as well as superior performance of I-TAP for all desirable levels of unfairness β , as compared to the algorithm in [16]. Moreover, our results demonstrate that our approach can reduce unfairness by 50% while increasing the total travel time by at most 2%, which indicates that a huge gain in user fairness can be achieved for a small loss in efficiency, making our approach a desirable option for use in route guidance systems.

This paper is organized as follows. Section 2 reviews related literature. We introduce in Section 3 the β -SO problem and metrics to evaluate the fairness and efficiency of a traffic assignment. We introduce the I-TAP method and discuss its properties in Section 4,

and develop pricing schemes in Section 5. We evaluate the performance of the I-TAP method through numerical experiments in Section 6 and provide directions for future work in Section 7.

In the extended version of our paper [17], we provide omitted proofs, numerical implementation details of our approach and extensions to fairness notions beyond the one considered here.

2 RELATED WORK

The trade-off between system efficiency and user fairness has been widely studied in applications including resource allocation, reducing the bias of machine-learning algorithms, and influence maximization. While different notions of fairness have been proposed, the level of fairness is typically controlled through the problem’s objective or constraints. For instance, fairness parameters that trade-off the level of fairness in the objective can be tuned to investigate the loss in system efficiency in the context of influence maximization [21] and resource allocation [6] problems. On the other hand, fairness parameters that bound the degree of allowable inequality between different user groups through the problem’s constraints have been proposed to reduce bias towards disadvantaged groups [28], e.g., through group-based or diversity constraints [15, 29].

In the context of traffic routing, several traffic assignment formulations have been proposed to achieve a balance between multiple performance criteria [9, 10], with a particular focus on fairness considerations in traffic routing [16]. Since Jahn et al. [16] introduced the CSO problem, there have been both theoretical studies [25] as well as the development of heuristic approaches to solve the NP-hard CSO problem. For instance, [16] proposed a Frank-Wolfe based heuristic, while others have considered linear relaxations of the original problem [1–3]. Each of these approaches bounds the level of unfairness in terms of normal lengths of paths by restricting the set of eligible paths on which users can travel to those that meet a specified level of normal unfairness. However, the experienced unfairness in terms of the travel times may be much higher than the level of normal unfairness, which is an *a priori* fixed quantity.

This inherent drawback of the CSO problem in limiting the experienced unfairness in terms of user travel times was overcome by [4], which proposed two Mixed Integer Non-Linear Programming models to capture traffic-dependent notions of unfairness. Their approach to solve these models relies on a linearization heuristic for the edge travel-time functions, which are in general *non-linear*. Achieving a high level of accuracy of the linear relaxations in approximating the true travel-time functions, however, requires solving a large MILP which is computationally *expensive*. Unlike [4], our I-TAP method is computationally *inexpensive*, while directly accounting for non-linear travel-time functions.

A further limitation of the existing methods for fair traffic routing is that there are limited results in providing pricing schemes to induce selfish users to collectively form the proposed traffic patterns, e.g., those satisfying a certain bound on unfairness. For instance, [11] provides tolling mechanisms to enforce fairness constrained flows which applies only to parallel networks. In more general networks, [19] proposes an auction-based bidding mechanism for users to be assigned to precomputed paths. However, this approach cannot be applied as-is to our setting as users are unconstrained with respect to a specific path set.

3 MODEL AND PROBLEM DEFINITION

We model the road network as a directed graph $G = (V, E)$, with the vertex and edge sets V and E , respectively. Each edge $e \in E$ has a normal length η_e , and a flow-dependent travel-time function $t_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{>0}$, which maps x_e , the rate of traffic on edge e , to the travel time $t_e(x_e)$. As is standard in the traffic routing literature, we assume that the function t_e , for each $e \in E$, is differentiable, convex, locally Lipschitz continuous, and monotonically increasing.

Users make trips between a set of O-D pairs, and we model users with the same origin and destination as one commodity, where K is the set of all commodities. Each commodity $k \in K$ has a demand rate $d_k > 0$, which represents the amount of flow to be routed on a set of directed paths $\mathcal{P}_k$ between its origin and destination. The edge flow of each commodity k is given by $\mathbf{x}^k = \{x_e^k\}_{e \in E}$, while the aggregate edge flow is denoted as $\mathbf{x} := \{x_e\}_{e \in E}$. For an edge flow $\mathbf{x} := \{x_e\}_{e \in E}$ and a path $P \in \mathcal{P} = \cup_{k \in K} \mathcal{P}_k$, the amount of flow routed on the path is denoted as $\mathbf{x}_P$, where the vector of path flows $\mathbf{f} = \{\mathbf{x}_P : P \in \mathcal{P}\}$. Then, the travel time on path P is $t_P(\mathbf{x}) = \sum_{e \in P} t_e(x_e)$, while $\eta_P = \sum_{e \in P} \eta_e$ is its normal length.

We assume users are selfish and thus choose paths that minimize their total travel cost that is a linear function of tolls and travel time. For a value of time parameter $v > 0$, and a vector of edge prices (or tolls) $\tau = \{\tau_e\}_{e \in E}$, the travel cost on a given path P under the traffic assignment $\mathbf{x}$ is given by $C_P(\mathbf{x}, \tau) = v t_P(\mathbf{x}) + \sum_{e \in P} \tau_e$.

3.1 Traffic Assignment

In this work we will consider several variants of the traffic assignment problem (TAP). The goal of the SO traffic assignment problem (SO-TAP) is to route users to minimize the total system travel time. This behavior is captured in the following convex program:

DEFINITION 1 (PROGRAM FOR SO-TAP [27]).

$$\min_{\mathbf{f}} \quad h^{SO}(\mathbf{x}) := \sum_{e \in E} x_e t_e(x_e), \quad (1a)$$

$$\text{s.t.} \quad \sum_{k \in K} \sum_{P \in \mathcal{P}_k: e \in P} \mathbf{x}_P = x_e, \quad \forall e \in E, \quad (1b)$$

$$\sum_{P \in \mathcal{P}_k} \mathbf{x}_P = d_k, \quad \forall k \in K, \quad (1c)$$

$$\mathbf{x}_P \geq 0, \quad \forall P \in \mathcal{P}, \quad (1d)$$

with edge flow Constraints (1b), demand Constraints (1c), and non-negativity Constraints (1d).

We mention that the total travel time objective is only a function of the aggregate edge flow $\mathbf{x}$, which is related to the path flow $\mathbf{f}$ through Constraint (1b). Note for any given path flow $\mathbf{f}$ that both the edge flow $\mathbf{x}$ and the commodity-specific edge flows $\mathbf{x}^k$ for each commodity $k \in K$ are uniquely defined. Closely related to SO-TAP is the UE traffic assignment problem (UE-TAP) that emerges from the selfish behavior of users that minimize their own travel time, and is described by the following convex program:

DEFINITION 2 (PROGRAM FOR UE-TAP [27]).

$$\min_{\mathbf{f}} \quad h^{UE}(\mathbf{x}) := \sum_{e \in E} \int_0^{x_e} t_e(y) dy, \quad (2a)$$

$$\text{s.t.} \quad (1b) - (1d). \quad (2b)$$

While the integral objective used to define UE-TAP has not found a clear economic or behavioral interpretation within the transportation and game-theory communities [27], the optimal solution of UE-TAP corresponds to an equilibrium, which can be seen through the KKT conditions of this optimization problem. That is, UE-TAP provides a polynomial time computable method to determine the user equilibrium flows. A defining property of the UE solution is that it is fair for all users since the travel time of all the flow that is routed between the same O-D pair is equal. In contrast, at the SO solution the sum of the travel time and marginal cost of travel is the same for all users travelling between the same O-D pair. Thus, marginal cost pricing is used to induce selfish users to collectively form the SO traffic pattern. While the number of constraints, which depend on the path sets $\mathcal{P}_k$, can be exponential in the size of the transportation network, both SO-TAP and UE-TAP are efficiently computable since they can be formulated without explicitly enumerating all the path level flows and constraints [27].

3.2 Fairness and Efficiency Metrics

We evaluate the quality of any traffic assignment $\mathbf{x}$ using two metrics, namely: (i) efficiency and (ii) fairness.

We evaluate the efficiency of a traffic assignment by comparing its total travel time to that of the SO edge flow $\mathbf{x}^{SO}$. Recalling that $h^{SO}(\mathbf{x})$ denotes the total travel time of the edge flow $\mathbf{x}$, the *inefficiency ratio* of $\mathbf{x}$ is $\rho(\mathbf{x}) := \frac{h^{SO}(\mathbf{x})}{h^{SO}(\mathbf{x}^{SO})}$. Note that for the UE solution $\mathbf{x}^{UE}$, the inefficiency ratio is the Price of Anarchy (PoA) [18].

To evaluate the fairness of a traffic assignment, we first introduce the notion of a *positive path* from [5].

DEFINITION 3 (POSITIVE PATH). For any path flow $\mathbf{f}$ with corresponding commodity-specific edge flows $\mathbf{x}^k$, a path $P \in \mathcal{P}_k$ is *positive* for a commodity $k \in K$ if for all edges $e \in P$, x_e^k is strictly positive. The set of all positive paths for a flow $\mathbf{f}$ and commodity k is denoted as $\mathcal{P}_k^+(\mathbf{f}) = \{P : P \in \mathcal{P}_k, x_e^k > 0, \text{ for all } e \in P\}$.

The importance of the notion of a positive path is that the path decomposition of the commodity-specific edge flows $\mathbf{x}^k$ may be non-unique; however the set of positive paths is always uniquely defined for such edge flows. That is, for commodity specific edge flows $\mathbf{x}^k$ the set of positive paths for any two path decompositions $\mathbf{f}_1$ and $\mathbf{f}_2$ are equal, i.e., $\mathcal{P}_k^+(\mathbf{f}_1) = \mathcal{P}_k^+(\mathbf{f}_2)$.

We evaluate the fairness of a traffic flow $\mathbf{f}$ with an edge decomposition $\mathbf{x}$ through its corresponding unfairness U , which is defined as the maximum ratio across all O-D pairs of (i) the travel time on the slowest, i.e., highest travel time, positive path to (ii) the travel time on the fastest positive path between the same O-D pair, i.e., $U(\mathbf{f}) := \max_{k \in K} \max_{Q, R \in \mathcal{P}_k^+} \frac{t_Q(\mathbf{x})}{t_R(\mathbf{x})}$. That is, $U(\mathbf{f})$ returns the maximum possible ratio of travel times on positive paths across all commodities with respect to the path flow $\mathbf{f}$. As a result, the unfairness U is a number between one and infinity, and a traffic assignment has a high level of fairness if its unfairness is close to one while it has a low level of fairness if the corresponding unfairness is much larger than one. In contrast, other valid notions of unfairness could also be considered. For instance, for a given path flow decomposition $\mathbf{f}$ with a corresponding edge flow $\mathbf{x}$, the unfairness $\tilde{U}(\cdot)$ of the path flows can be evaluated as the maximum

ratio between the travel times of any two users travelling between the same O-D pair, i.e., $\tilde{U}(\mathbf{f}) = \max_{k \in K} \max_{Q, R \in \mathcal{P}_k: \mathbf{x}_Q, \mathbf{x}_R > 0} \frac{t_Q(\mathbf{x})}{t_R(\mathbf{x})}$. Note here that we only consider a ratio of travel times on paths with strictly positive flow for the path decomposition $\mathbf{f}$ rather than the ratio of travel times for all positive paths. We defer a detailed treatment of other path-based unfairness measures to the extended version of our paper [17] and highlight here some key features of the positive path based unfairness measure U .

The unfairness measure $U(\mathbf{f})$ can be efficiently computed and has the benefit that it applies to all possible path decompositions of the commodity specific edge flows $\mathbf{x}^k$. As a result, in the context of a single O-D pair travel demand, the unfairness measure $U(\mathbf{f})$ has the benefit that it is a property of the unique edge flow $\mathbf{x}$ and is relevant when users are not constrained to a specific path decomposition, as happens in practice. In contrast, path decomposition specific unfairness measures, e.g., $\tilde{U}(\mathbf{f})$, are likely to be more sensitive to the method used to compute the path decomposition. Furthermore, we note that the positive path based unfairness notion serves as an upper bound on the ratio of travel times for any two users travelling between the same O-D pair for the path flow $\mathbf{f}$, i.e., $\tilde{U}(\mathbf{f}) \leq U(\mathbf{f})$ for all $\mathbf{f}$. As a result, our theoretical bounds on unfairness obtained for the positive path-based unfairness notion will naturally extend to path decomposition specific unfairness measures such as $\tilde{U}(\mathbf{f})$. Thus, in the rest of this paper we focus on the positive path-based unfairness measure and, for numerical comparison, we present other path decomposition specific unfairness measures, e.g., $\tilde{U}(\mathbf{f})$, in the extended version of our paper [17].

3.3 Toy Network Example

To illustrate the fairness and efficiency properties of the two optimization problems, SO-TAP and UE-TAP, we present a toy example of a two-edge Pigou network, as depicted in Figure 1. In particular, consider a demand of one that needs to be routed from the origin v_1 to the destination v_2 , with two edges (e_1 and e_2) connecting the origin to the destination. Observe that if the travel time functions on the two edges are given by $t_1(x_1) = 1$ and $t_2(x_2) = x_2$, then under the UE-TAP solution all users will be routed on edge two, while the SO-TAP solution that minimizes the total travel time will route 0.5 units of flow on both edges. The level of unfairness and the total travel time of the two traffic assignments are presented in the following table, which indicates that the UE-TAP solution is fair but inefficient while the SO-TAP solution is efficient but unfair.




Figure 1: A two-edge Pigou network to illustrate the fairness and efficiency properties of SO-TAP and UE-TAP. For a demand of one, the UE-TAP solution routes all the flow on edge e_2 resulting in a fair solution but a total travel time of one. On the other hand, the SO-TAP solution routes 0.5 units of flow on both edges, resulting in an efficient solution with the minimum total travel time but an unfairness of two.

3.4 β -Fair System Optimum

To trade-off between user fairness and system efficiency, we consider the following β -fair System Optimum (β -SO) problem, wherein we impose an upper bound $\beta \in [1, \infty)$ on the maximum allowable unfairness in the network, i.e., $U(\mathbf{f}) \leq \beta$ for a path flow $\mathbf{f}$.

DEFINITION 4 (PROGRAM FOR β -FAIR SYSTEM OPTIMUM).

$$\min_{\mathbf{f}} \sum_{e \in E} x_e t_e(x_e), \quad (3a)$$

$$\text{s.t.} \quad (1b) - (1d), \quad (3b)$$

$$U(\mathbf{f}) \leq \beta. \quad (3c)$$

Note that without the unfairness Constraints (3c) (or when $\beta = \infty$), the above problem exactly coincides with SO-TAP. Furthermore, the β -SO problem is always feasible for any $\beta \in [1, \infty)$, since a solution to UE-TAP exists and achieves an unfairness of $\beta = 1$.

We also note that the difference between the β -SO and CSO problems is in the unfairness Constraints (3c). While the β -SO problem explicitly imposes an upper limit on the ratio of travel times on positive paths, the CSO problem imposes normal unfairness constraints for each path $P \in \mathcal{P}_k$ and any commodity $k \in K$ of the form $\eta_P \leq \phi \min_{P^* \in \mathcal{P}_k} \eta_{P^*}$ for some normal unfairness parameter $\phi \geq 1$. That is, the CSO problem minimizes the total travel time subject to flow conservation constraints over the set of paths with a normal unfairness level of at most ϕ . The authors of [16] use normal unfairness, which is a fixed quantity, as a proxy to limit the ratio of user travel times, which vary according to a traffic assignment.

The optimal solution of the β -SO problem corresponds to the highest achievable system efficiency whilst meeting unfairness constraints. However, solving β -SO directly is generally intractable as the unfairness Constraints (3c) are non-convex if the travel time function is non-linear. Moreover, since the unfairness metric studied in this work accounts for user costs that vary according to a traffic assignment, unlike normal unfairness that is an apriori fixed quantity, the NP-hardness of the CSO problem [16] suggests the computational hardness of β -SO [4, 5].

Finally, we mention that we consider a setting wherein the travel demand is time invariant and fractional user flows are allowed, as is standard in the traffic routing literature. Also, for notational simplicity, we consider for now a model where all users are homogeneous, i.e., they have an identical value of time v , and present an extension of our pricing result to the setting of heterogeneous users in Section 5.2.

4 A METHOD FOR β -FAIR SYSTEM OPTIMUM

In this section, we develop a computationally-efficient method for solving β -SO with edge-based unfairness constraints, to achieve a traffic assignment with a low total travel time, whose level of unfairness is at most β . In particular, we propose a new formulation of TAP, which we term interpolated TAP (or I-TAP), wherein the objective function linearly interpolates between the objectives of UE-TAP and SO-TAP. Our main insight is that the UE solution achieves a high level of fairness, whereas the SO solution achieves a low total travel time, and we wish to get the best of both worlds—high level of fairness at a low total travel time.

In this section, we describe the I-TAP method and evaluate its efficacy for the β -SO problem by addressing three key concerns regarding the solution efficiency, feasibility and computational tractability. In particular, we establish a relationship between I-TAP and β -SO through theoretical bounds on the inefficiency ratio (Section 4.2) and its optimality for two-edge Pigou networks (Section 4.3). We also establish the feasibility of I-TAP for β -SO by finding the range of values of the interpolation parameter such that the unfairness of the optimal I-TAP solution is guaranteed to be less than β (Section 4.2). Finally, we present an equivalence between I-TAP and UE-TAP to show that I-TAP can be computed efficiently (Section 4.4). These results indicate that we can approximate β -SO as an unconstrained traffic assignment problem that can be solved quickly. We also mention that we perform a sensitivity analysis to establish the continuity of the optimal traffic assignment and its total travel time in the interpolation parameter of I-TAP in the extended version of our paper [17].

4.1 Interpolated Traffic Assignment

We provide a formal definition of interpolated TAP:

DEFINITION 5 (I-TAP). *For a convex combination parameter $\alpha \in [0, 1]$, the interpolated traffic assignment problem, denoted as I-TAP $_{\alpha}$, is given by:*

$$\min_{\mathbf{f}} \quad h^I_{\alpha}(\mathbf{x}) := \alpha h^{SO}(\mathbf{x}) + (1 - \alpha) h^{UE}(\mathbf{x}), \quad (4a)$$

$$\text{s.t.} \quad (1b) - (1d). \quad (4b)$$

A few comments are in order. First, it is clear that I-TAP $_0$ and I-TAP $_1$ correspond to UE-TAP and SO-TAP, respectively. Next, under the assumption that the travel time functions are strictly convex, we observe that for any $\alpha \in [0, 1]$ the program I-TAP $_{\alpha}$ is a convex optimization problem with a unique edge flow solution $\mathbf{x}(\alpha)$.

For numerical implementation purposes, we propose a *dense sampling* procedure to compute a solution for β -SO with a low total travel time while guaranteeing a β bound on unfairness. In particular, to compute a good solution for β -SO, we evaluate the optimal solution $\mathbf{f}(\alpha)$ of I-TAP $_{\alpha}$ (with corresponding edge flows $\mathbf{x}(\alpha)$) for α taken from a finite set $\mathcal{A}_s := \{0, s, 2s, \dots, 1\}$ for some step size $s \in (0, 1)$. That is, in $O(\frac{1}{s})$ computations of I-TAP, we can return the path flow $\mathbf{f}(\alpha^*)$ (with edge decomposition $\mathbf{x}(\alpha^*)$), for some $\alpha^* \in \mathcal{A}_s$, with the lowest total travel time that is at most β -unfair, i.e., $U(\mathbf{f}(\alpha^*)) \leq \beta$, and the value $h^{SO}(\mathbf{x}(\alpha^*))$ is minimized.

We observe experimentally (Section 6.2) that this method of computing the I-TAP solution achieves a good solution for β -SO in terms of fairness and total travel time. We note here that our approach also naturally extends to other unfairness notions wherein the user equilibrium achieves the highest possible level of fairness, while the system optimum achieves the lowest total travel times (see the extended version of our paper [17]). Finally, we restrict α to lie in the finite set $\mathcal{A}_s$ since the exact functional form of the optimal solution $\mathbf{f}(\alpha)$ (with edge flow $\mathbf{x}(\alpha)$), and thus the unfairness $U(\mathbf{f}(\alpha))$ and the total travel time $h^{SO}(\mathbf{f}(\alpha))$ functions, in α is not directly known, though we show that $\mathbf{x}(\alpha)$ and $h^{SO}(\mathbf{x}(\alpha))$ are continuous in α in the extended version of our paper [17].

We also test (Section 6) an alternative approach to I-TAP, which instead of taking a convex combination of the SO-TAP and UE-TAP

objectives, interpolates between their unique edge flow solutions. That is, we first compute the optimal edge flow solutions of UE-TAP ($\mathbf{x}^{UE}$) and SO-TAP ($\mathbf{x}^{SO}$), and return the value $(1 - \gamma)\mathbf{x}^{UE} + \gamma\mathbf{x}^{SO}$ for $\gamma \in [0, 1]$. While this Interpolated Solution (I-Solution) method only requires two traffic assignment computations as compared to $O(\frac{1}{s})$ computations of the I-TAP method, it leads to poor performance in comparison with the I-TAP method (see Section 6) and does not induce a natural marginal-cost pricing scheme, as I-TAP does (see Section 5). Thus, we focus on I-TAP for this and the next sections.

4.2 Solution Efficiency and Fairness of I-TAP

In this section, we study the influence of the convex combination parameter α of I-TAP on the efficiency and fairness of the optimal solution $\mathbf{f}(\alpha)$ (and edge flow $\mathbf{x}(\alpha)$). In particular, we characterize (i) an upper bound on the inefficiency ratio as we vary α , and (ii) a range of values of α that are guaranteed to achieve a specified level of unfairness β for any optimal solution $\mathbf{f}(\alpha)$.

We first evaluate the performance of I-TAP by establishing an upper bound on the inefficiency ratio of the optimal solution of I-TAP $_{\alpha}$ as a function of α . Theorem 1 shows that this bound is a minimum between (i) the PoA, which we denote as $\bar{\rho}$, and (ii) a more elaborate bound that is monotonically non-increasing in α .

THEOREM 1 (I-TAP SOLUTION EFFICIENCY). *For any $\alpha \in (0, 1)$, let $\mathbf{x}(\alpha)$ be the optimal edge flow of I-TAP $_{\alpha}$. Then, the inefficiency ratio*

$$\rho(\mathbf{x}(\alpha)) \leq \min \left\{ \bar{\rho}, 1 + \frac{1 - \alpha}{\alpha} \cdot \frac{h^{UE}(\mathbf{x}(1)) - h^{UE}(\mathbf{x}(0))}{h^{SO}(\mathbf{x}(1))} \right\}.$$

For a proof of Theorem 1 and all subsequent results, see the extended version of our paper [17]. Theorem 1 establishes that, even in the worst case, the ratio between the total travel time of the edge flow $\mathbf{x}(\alpha)$ and that of the system optimal solution is at most the PoA. This result is not guaranteed to hold for other state-of-the-art CSO algorithms, e.g., the algorithm in [16] (see Section 6). Further, Theorem 1 shows that the upper bound on the inefficiency ratio becomes closer to one as the objective h^I_{α} gets closer to h^{SO} .

We now establish a range of values of α at which the any optimal solution $\mathbf{f}(\alpha)$ of I-TAP $_{\alpha}$ is guaranteed to attain a β bound on unfairness for polynomial travel time functions, e.g., the commonly used BPR function [27].

THEOREM 2 (FEASIBILITY OF I-TAP FOR β -SO). *Suppose that the largest degree of the polynomial travel time functions $t_e(x_e)$ is m for some $e \in E$. Then, the unfairness of any optimal solution $\mathbf{f}(\alpha)$ of I-TAP $_{\alpha}$ is upper bounded by β , i.e., $U(\mathbf{f}(\alpha)) \leq \beta$, for any $\alpha \leq \frac{\beta-1}{m}$.*

We can further show that the bound in Theorem 2 is in fact tight by demonstrating an instance such that for any $\alpha > \frac{\beta-1}{m}$ the unfairness of the solution $\mathbf{f}(\alpha)$ of I-TAP $_{\alpha}$ is strictly greater than β .

LEMMA 1 (TIGHTNESS OF UNFAIRNESS BOUND). *Suppose $\mathbf{f}(\alpha)$ is an optimal solution to I-TAP $_{\alpha}$ for any $\alpha \in [0, 1]$. Then, there exists a two-edge parallel network with polynomial travel time functions of degree at most m such that for any $\alpha > \frac{\beta-1}{m}$, the unfairness $U(\mathbf{f}(\alpha)) > \beta$.*

Together, Theorem 2 and Lemma 1 imply that a β level of unfairness can be guaranteed using I-TAP on all traffic networks only when $\alpha \leq \frac{\beta-1}{m}$, where m is the maximum degree of the polynomial corresponding to the travel time functions for each edge $e \in E$.

4.3 Optimality of I-TAP

In this section, we show that I-TAP exactly computes the minimum total travel time solution for any desired level of unfairness β in any two edge Pigou network. That is, there is some convex combination parameter α^* for which the solution of I-TAP $_{\alpha^*}$ is also a solution to the β -SO problem for any two edge Pigou network.

LEMMA 2 (OPTIMALITY OF I-TAP). *Consider a two edge Pigou network where the optimal solution of the β -SO problem is $\mathbf{x}_{\beta}^*$ for any $\beta \in [1, \infty)$. Then, there exists a convex combination parameter α^* such that $\mathbf{x}(\alpha^*) = \mathbf{x}_{\beta}^*$.*

We mention that Lemma 2 compares only the edge flows of I-TAP and β -SO since the path and edge flows coincide for a two edge Pigou network. We also note that while the optimality for a Pigou network may appear restrictive, such networks are of both theoretical [20, 24] and practical significance [8].

4.4 Computational Tractability of I-TAP

Having established that we can solve I-TAP to obtain an approximate solution to β -SO, we now establish that I-TAP can be computed efficiently due to its equivalence to a parametric UE-TAP program.

OBSERVATION 1 (UE EQUIVALENCY OF I-TAP). *For any $\alpha \in [0, 1]$, I-TAP $_{\alpha}$ reduces to UE-TAP with objective function $\sum_{e \in E} \int_0^{x_e} c_e(y, \alpha) dy$, where $c_e(y, \alpha) = t_e(y) + \alpha y t'_e(y)$.*

Observation 1 follows from the fundamental theorem of calculus. Note that for each $\alpha \in [0, 1]$, the differentiability, monotonicity, and convexity of many typical travel time functions t_e , e.g., any polynomial function such as the BPR function [27], imply that the corresponding properties hold for the cost functions $c_e(x_e, \alpha)$ in x_e . For numerical implementation, the equivalency of I-TAP $_{\alpha}$ and UE-TAP implies that I-TAP $_{\alpha}$ inherits the useful property that the linearization step of the Frank-Wolfe algorithm [27], when applied to I-TAP $_{\alpha}$, corresponds to solving multiple unconstrained shortest path queries. The latter motivates the highly efficient approach which we employ in Section 6 to solve I-TAP $_{\alpha}$.

5 PRICING TO IMPLEMENT FLOWS

In this section, we leverage the structure of I-TAP to develop pricing mechanisms to collectively enforce the I-TAP solution in the presence of selfish users that independently choose routes to minimize their own travel costs. We first consider the case of homogeneous users and show that I-TAP results in a natural marginal-cost pricing scheme. Then, we characterize conditions under which tolls can be used to enforce the I-TAP flows for heterogeneous users.

In this section, for the ease of exposition, we focus our discussion on inducing the optimal edge flow $\mathbf{x}(\alpha)$ of I-TAP $_{\alpha}$. We mention that our approach can naturally be extended to enforcing optimal path flows $\mathbf{f}(\alpha)$ that satisfy a given level of unfairness. In particular, we can consider a setting wherein users are recommended to use a specified path set, e.g., by traffic navigational applications, as given by $\mathbf{f}(\alpha)$ and the tolls set are such that no user will have an incentive to deviate from their recommended paths. Finally, we also mention by Theorem 2 that focusing on the edge flow $\mathbf{x}(\alpha)$ is without loss of generality for certain ranges of α since the unfairness bound for any optimal path flow solution $\mathbf{f}(\alpha)$ is guaranteed to be satisfied.

5.1 Homogeneous Pricing via Marginal Cost

In the setting where all users have the same value of time v , the structure of I-TAP $_{\alpha}$ yields an interpolated variant of marginal-cost pricing to induce selfish users to collectively form the optimal edge flow $\mathbf{x}(\alpha)$ of I-TAP $_{\alpha}$. This result is a direct consequence of the equivalence between I-TAP and UE-TAP from Observation 1.

LEMMA 3 (PRICES TO IMPLEMENT FLOWS). *Suppose that $\mathbf{x}(\alpha)$ is a solution to I-TAP $_{\alpha}$ for some $\alpha \in [0, 1]$. Then $\mathbf{x}(\alpha)$ can be enforced as a UE by setting the prices as $\tau_e = \alpha x_e(\alpha) t'_e(x_e(\alpha))$ for each $e \in E$.*

Note from Lemma 3 that the edge prices are equal to α multiplied by the marginal cost of users.

5.2 Heterogeneous Pricing via Dual Multipliers

The pricing mechanism in Section 5.1 is inapplicable to the heterogeneous user setting as it would require unrealistically imposing different prices for users with different values of time for the same edges. In this section, we consider heterogeneous users and leverage a linear-programming method [12] to establish that appropriate tolls can be placed on the roads to induce heterogeneous selfish users to collectively form the equilibrium edge flow $\mathbf{x}(\alpha)$.

Before presenting the pricing scheme, we first extend the notion of a commodity to a heterogeneous user setting. In particular, each user belongs to a commodity $k \in K$ when making a trip on a set of paths $\mathcal{P}_k$ between the same O-D pair and has the value of time $v_k > 0$. Then, under a vector of edge prices $\tau = \{\tau_e\}_{e \in E}$ the travel cost that users in commodity k incur on a given path $P \in \mathcal{P}_k$ under the traffic assignment $\mathbf{x}$ is given by $C_P(\mathbf{x}, \tau) = \sum_{e \in P} (v_k t_e(x_e) + \tau_e)$. Note that more than one commodity may make trips between the same O-D pair, and a user equilibrium forms when the travel cost for all users in a particular commodity is equal. We further note that we maintain the unfairness notion presented in the work even for heterogeneous users. That is, irrespective of the value of time of two users travelling between the same O-D pair, the maximum possible ratio between their travel times can be no more than β .

We now leverage the following result to provide a necessary and sufficient condition that the optimal edge flow $\mathbf{x}(\alpha)$ of I-TAP $_{\alpha}$ must satisfy for it to be enforceable as a UE through road pricing.

LEMMA 4 (CONDITION FOR FLOW ENFORCEABILITY). [12, Theorem 3.1] *Suppose that the non-negative flow $\mathbf{x}$ satisfies the edge flow and demand constraints in Definition 1. Further, consider the linear program: $\min_{d_P^k \in \tilde{\Omega}} \sum_{k \in K} v_k \sum_{P \in \mathcal{P}_k} t_P(\mathbf{x}) d_P^k$, where the non-negative variables d_P^k represent the flow of commodity k on path $P \in \mathcal{P}_k$, and $\mathcal{P}_k$ denotes the set of all possible paths for commodity k . Here $\tilde{\Omega}$ is the set described by non-negative flows satisfying capacity constraints, i.e., $\sum_{k \in K} \sum_{P \in \mathcal{P}_k: e \in P} d_P^k \leq x_e$ for all edges $e \in E$, and demand constraints, i.e., $\sum_{P \in \mathcal{P}_k} d_P^k = d_k$ for all $k \in K$. Then $\mathbf{x}$ can be enforced as a UE if and only if the capacity constraints are met with equality for each edge at the optimal solution of the linear program.*

We now show that $\mathbf{x}(\alpha)$ satisfies the condition in Lemma 4.

LEMMA 5 (HETEROGENEOUS USER FLOW ENFORCEABILITY). *Suppose that the edge flow $\mathbf{x}(\alpha)$ is a solution for I-TAP $_{\alpha}$ for some $\alpha \in [0, 1]$. Then for the heterogeneous user setting, $\mathbf{x}(\alpha)$ can be enforced as a user equilibrium.*

Table 1: Problem instance attributes and computation time. For each instance we report the number of vertices $|V|$, edges $|E|$, and OD pairs $|K|$. In addition, we report the computation time of each instance for the previous method of Jahn et al. [16] and our I-TAP method using 100 iterations of the Frank-Wolfe algorithm.

Region Name	attributes			runtime (sec.)	
	$ V $	$ E $	$ K $	Jahn et al.	I-TAP
Sioux Falls (SF)	24	76	528	20.0	0.03
Anaheim (A)	416	914	1406	74.0	0.33
Massachusetts (M)	74	258	1113	24.3	0.09
Tiergarten (T)	361	766	644	18.2	0.20
Friedrichshain (F)	224	523	506	19.8	0.12
Prenzlauerberg (P)	352	749	1406	74.4	0.32

Lemma 5 implies that even when users are heterogeneous the edge flow $\mathbf{x}(\alpha)$ can be enforced as an equilibrium flow using tolls set through the dual variables of a linear program.

6 NUMERICAL EXPERIMENTS

We now evaluate the performance of our I-TAP method for β -SO on several real-world transportation networks. The results of our experiments not only characterize the behavior of I-TAP but also highlight that, compared to the algorithm in [16], our approach has much smaller runtimes while achieving lower total travel times for most levels β of unfairness. We present the implementation details of the I-TAP method and the unfairness metric in the extended version of our paper [17]. In the following, we describe the data-sets we use and present the corresponding results to evaluate the performance of our approach.

6.1 Data Sets

Table 1 shows the six instances we use for our study, which were obtained from [14]. We use the BPR travel time function [27], defined as $t_e(x_e) = \xi_e(1 + 0.15(\frac{x_e}{\kappa_e})^4)$, where ξ_e is the free-flow travel time on edge e , and κ_e is the capacity of edge e , which is the number of users beyond which the travel time on the edge rapidly increases.

6.2 Results

Assessment of Theoretical Upper Bounds. We now assess the theoretical upper bounds on the inefficiency ratio and unfairness that were obtained in Section 4.2. We present the results for the Prenzlauerberg data-set and note that the results extend to other problem instances in Table 1 as well.

Figure 2 depicts both (i) the change in the inefficiency ratio (left) and unfairness (right) of the solution of I-TAP using dense sampling, and (ii) the theoretical upper bound of the inefficiency ratio (Theorem 1) and unfairness (Theorem 2). As expected, the dense sampling procedure results in both an inefficiency ratio and unfairness that is below the theoretical upper bound for every value of α .

Behavior of I-TAP. For each of the transportation networks in Table 1, we now study the relationship between the convex combination parameter α and the (i) total travel time, and (ii) unfairness.

Figure 3 (left) shows the relationship between the inefficiency ratio and α . Note that when $\alpha = 1$, the inefficiency ratio is one,




Figure 2: Comparison between the inefficiency ratio (left) and unfairness (right) of the solution of I-TAP sampling on the Prenzlauerberg data-set and the theoretical bounds obtained in Theorems 1 and 2. The convex combination parameters were chosen at increments of 0.01.

since the interpolated objective is the SO-TAP objective, and when $\alpha = 0$, the inefficiency ratio is the Price of Anarchy (PoA), since the interpolated objective is the UE-TAP objective.

The relationship between unfairness and α is depicted on the right in Figure 3, where for readability we have marked outliers where large changes in the unfairness occur for small changes in α . For an explanation of the jumps in the unfairness at certain values of α , see the extended version of our paper [17].

Finally, for each transportation network the general trend of a decrease in the inefficiency ratio and an increase in the unfairness with an increase in α suggests that decreasing the total travel time comes at the cost of an increase in the unfairness and vice versa.

Solution Quality Comparison. We now explore the efficiency-fairness tradeoff through a comparison of the Pareto frontier of the I-TAP method to the approach in [16], which is a benchmark solution for fair traffic routing, and the I-Solution method described in Section 4.1. To this end, we depict the Pareto frontier of the (i) I-TAP method for 0.01 and 0.05 increments of the parameter α , (ii) I-Solution method for 0.01 increments of the convex combination parameter γ , and (iii) Jahn et al.'s approach [16] for 0.05 increments of the normal unfairness parameter ϕ lying between one and two.

Figure 4 depicts the Pareto frontiers, i.e., the set of all Pareto efficient combinations of system efficiency and user fairness, for the six transportation networks in Table 1. In particular, observe that the Pareto frontiers of the I-TAP method are below that of the other two approaches for most values of unfairness. This observation indicates that the I-TAP method outperforms the other two approaches since the inefficiency ratio of the I-TAP solution is the lowest for most desired levels of unfairness. Only for the Sioux Falls and Prenzlauerberg data-sets, the algorithm in [16] achieved lower inefficiency ratios than both the I-TAP and I-Solution methods for higher values of unfairness, which, in practice, would be undesirable. Furthermore, note that, unlike the two convex-combination approaches, the solution of the algorithm in [16] can result in inefficiency ratios that are much greater than the PoA for low levels of unfairness. The I-TAP method outperforms the I-Solution method since the set of paths that users can traverse is not restricted to the union of the routes under the UE and SO solutions as is the case for the I-Solution method. In particular, there may be traffic assignments with lower total travel times that use paths not encapsulated by the restricted set of paths corresponding to the I-Solution method. Furthermore, while the PoA for each of the data-sets is quite low, some real-world transportation networks may have much




Figure 3: Variation in the inefficiency ratio (left) and the level of unfairness (right) of the optimal solution of $I-TAP_\alpha$ with the parameter $\alpha \in [0, 1]$ for six different transportation networks from Table 1. The values of the convex combination parameter were chosen at increments of 0.01.




Figure 4: Pareto frontier depicting the trade-off between efficiency and fairness for the (i) I-TAP method with a step size $s = 0.01$, (ii) I-TAP method with $s = 0.05$, (iii) I-Solution method with $s = 0.01$, and (iv) Jahn et al.'s method [16] with $s = 0.05$.

higher PoA values (even as high as two) [32], which would make the trade-off between efficiency and fairness even more prominent.

Runtime Comparison. We report in Table 1 the runtime of the Jahn et al. method [16] and our I-TAP method. For each instance we report the average runtime over the parameters ϕ and α for the competitor and our method, respectively. We observe that our approach is faster by at least three orders of magnitude. This is unsurprising since our method solves *unconstrained* shortest-path queries, which can be implemented in $O(|E| + |V| \log |V|)$ time, within each Frank-Wolfe iteration, whereas [16] solves *constrained* shortest-path queries which are NP-hard. We do mention that a more efficient implementation of constrained shortest-path query can be achieved by directly implementing a label-correcting algorithm rather than using the `r_c_shortest_paths` routine from Boost, which is overly general for our setting and hence less efficient. Nevertheless, even with this improvement it would still be much slower than the unconstrained near-linear algorithm. Notice that both approaches can be sped up via parallel computation of shortest-path queries, and our method can be made even faster through modern heuristics for shortest-path queries [7].

7 CONCLUSION AND FUTURE WORK

In this paper, we developed (i) a computationally efficient method for traffic routing that trades-off system efficiency and user fairness, and (ii) pricing schemes to enforce fair traffic assignments as a UE.

There are various directions for future research. First, it would be valuable to develop theoretical bounds for I-TAP to demonstrate its applicability to other notions of unfairness, some of which are studied in the extended version of our paper [17]. Next, it would be useful to investigate fairness notions that compare user travel times across O-D pairs. Finally, it would be interesting to study I-TAP's generalizability when accounting for costs beyond the travel times of users, such as environmental pollution and user discomfort.

8 ACKNOWLEDGEMENTS

We thank Kaidi Yang and anonymous reviewers for their detailed inputs and insightful feedback on an earlier version of this work. This work was supported in part by the National Science Foundation (NSF) CAREER Award CMMI1454737, NSF Award 1830554, Toyota Research Institute, the Israeli Ministry of Science and Technology grants no. 3-16079 and 3-17385, and the United States-Israel Binational Science Foundation grants no. 2019703 and 2021643.

REFERENCES

- [1] Enrico Angelelli, Idil Arsik, Valentina Morandi, Martin Savelsbergh, and Maria Grazia Speranza. 2016. Proactive route guidance to avoid congestion. *Transportation Research Part B: Methodological* 94 (2016), 1–21. <https://doi.org/10.1016/j.trb.2016.08.015>
- [2] Enrico Angelelli, Valentina Morandi, Martin Savelsbergh, and Maria Grazia Speranza. 2020. System optimal routing of traffic flows with user constraints using linear programming. *European Journal of Operational Research* (2020). <https://doi.org/10.1016/j.ejor.2020.12.043>
- [3] Enrico Angelelli, Valentina Morandi, and Maria Grazia Speranza. 2018. Congestion avoiding heuristic path generation for the proactive route guidance. *Computers & Operations Research* 99 (2018), 234–248. <https://doi.org/10.1016/j.cor.2018.07.009>
- [4] Enrico Angelelli, Valentina Morandi, and Maria Grazia Speranza. 2020. Minimizing the total travel time with limited unfairness in traffic networks. *Computers & Operations Research* 123 (2020), 105016. <https://doi.org/10.1016/j.cor.2020.105016>
- [5] Soumya Basu, Ger Yang, Thanasis Lianes, Evdokia Nikolova, and Yitao Chen. 2017. Reconciling Selfish Routing with Social Good. *CoRR* abs/1707.00208 (2017).
- [6] Dimitris Bertsimas, Vivek Farias, and Nikolaos Trichakis. 2012. On the Efficiency-Fairness Trade-off. *Management Science* 58, 12 (2012), 2234–2250. <https://doi.org/10.1287/mnsc.1120.1549>
- [7] Valentin Buchhold, Peter Sanders, and Dorothea Wagner. 2019. Real-Time Traffic Assignment Using Engineered Customizable Contraction Hierarchies. *ACM J. Exp. Algorithmics* 24, Article 2.4 (Dec. 2019), 28 pages. <https://doi.org/10.1145/3362693>
- [8] Caltrans. 2010. US 101 South, corridor system management plan, 2010. Accessed July 01, 2020.
- [9] Robert Dial. 1996. Bicriterion Traffic Assignment: Basic Theory and Elementary Algorithms. *Transportation Science* 30, 2 (1996), 93–111. <https://doi.org/10.1287/trsc.30.2.93>
- [10] Robert Dial. 1997. Bicriterion traffic assignment: Efficient algorithms plus examples. *Transportation Research Part B: Methodological* 31, 5 (1997), 357–379. [https://doi.org/10.1016/S0191-2615\(96\)00034-3](https://doi.org/10.1016/S0191-2615(96)00034-3)
- [11] Bryce Ferguson and Jason Marden. 2021. The Impact of Fairness on Performance in Congestion Networks. In *2021 American Control Conference (ACC)*. 4521–4526. <https://doi.org/10.23919/ACC50511.2021.9483197>
- [12] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *IEEE Symp. on Foundations of Computer Science*. <https://doi.org/10.1109/FOCS.2004.69>
- [13] Sean Fleming. 2019. Traffic congestion cost the US economy nearly \$87 billion in 2018. <https://www.weforum.org/agenda/2019/03/traffic-congestion-cost-the-us-economy-nearly-87-billion-in-2018/> Accessed February 19, 2021.
- [14] Transportation Networks for Research Core Team. 2016. Transportation Networks for Research. github.com/bstabler/TransportationNetworks. Accessed January 20, 2021.
- [15] Lily Hu and Yiling Chen. 2020. Fair Classification and Social Welfare. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 535–545. <https://doi.org/10.1145/3351095.3372857>
- [16] O. Jahn, R. Möhring, A. Schulz, and N. Stier-Moses. 2005. System-Optimal Routing of Traffic Flows with User Constraints in Networks with Congestion. *Operations Research* 53, 4 (2005), 600–616. <https://doi.org/10.1287/opre.1040.0197>
- [17] Devansh Jalota, Kiril Solovey, Matthew Tsao, Stephen Zoepf, and Marco Pavone. 2021. Balancing Fairness and Efficiency in Traffic Routing via Interpolated Traffic Assignment. *arXiv:2104.00098 [eess.SY]*
- [18] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-Case Equilibria. In *Theoretical Aspects of Computer Science*.
- [19] Marin Lujak, Stefano Giordani, and S. Ossowski. 2014. Route Guidance: Bridging System and User Optimization in Traffic Assignment. *Neurocomputing* 151 (09 2014). <https://doi.org/10.1016/j.neucom.2014.08.071>
- [20] Arthur Pigou. 1912. *Wealth and Welfare* (1 ed.). London, Macmillan and Co.
- [21] Aida Rahmattalabi, Shahin Jabbari, Himabindu Lakkaraju, Phebe Vayanos, E. Rice, and Milind Tambe. 2020. Fair Influence Maximization: A Welfare Optimization Approach. *CoRR* abs/2006.07906 (2020).
- [22] Tim Roughgarden. 2002. How unfair is optimal routing?. In *ACM-SIAM Symp. on Discrete Algorithms*.
- [23] Tim Roughgarden and Éva Tardos. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47, 2 (2004), 389–403. <https://doi.org/10.1016/j.geb.2003.06.004>
- [24] Tim Roughgarden and Éva Tardos. 2002. How Bad is Selfish Routing? *J. ACM* 49, 2 (March 2002), 236–259. <https://doi.org/10.1145/506147.506153>
- [25] Andreas Schulz and Nicolás Stier-Moses. 2006. Efficiency and fairness of system-optimal routing with user constraints. *Networks* 48, 4 (2006), 223–234. <https://doi.org/10.1002/net.20133>
- [26] Guni Sharon, Michael Albert, Tarun Rambha, Stephen Boyles, and Peter Stone. 2018. Traffic Optimization For a Mixture of Self-interested and Compliant Agents. In *ISAIM*.
- [27] Yossi Sheffi. 1985. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods* (1 ed.). Prentice-Hall, Englewood Cliffs, New Jersey.
- [28] Ana-Andreea Stoica, Jessy Xinyi Han, and Augustin Chaintreau. 2020. Seeding Network Influence in Biased Networks and the Benefits of Diversity. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 2089–2098. <https://doi.org/10.1145/3366423.3380275>
- [29] Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, and Yair Zick. 2019. Group-Fairness in Influence Maximization. *CoRR* abs/1903.00967 (2019).
- [30] John Glen Wardrop. 1952. Some Theoretical Aspects of Road Traffic Research. *Proceedings of the Institution of Civil Engineers* 1, 3 (1952), 325–362. <https://doi.org/10.1680/ipeds.1952.11259>
- [31] Hai Yang and Hai-Jun Huang. 2005. *Mathematical and Economic Theory of Road Pricing* (1 ed.). Emerald Publishing.
- [32] Jing Zhang, Sepideh Pourazarm, Christos Cassandras, and Ioannis Paschalidis. 2018. The Price of Anarchy in Transportation Networks: Data-Driven Evaluation and Reduction Strategies. *Proc. IEEE* 106, 4 (2018), 538–553. <https://doi.org/10.1109/JPROC.2018.2790405>