

Average Unfairness in Routing Games

Pan-Yang Su*
University of California, Berkeley
Berkeley, United States
pan_yang_su@berkeley.edu

Arwa Alanqary*
University of California, Berkeley
Berkeley, United States
arwa@berkeley.edu

Bryce L. Ferguson
Dartmouth College
Hanover, United States
pan_yang_su@berkeley.edu

Manxi Wu
University of California, Berkeley
Berkeley, United States
manxiwu@berkeley.edu

Alexandre M. Bayen
University of California, Berkeley
Berkeley, United States
bayen@berkeley.edu

Shankar Sastry
University of California, Berkeley
Berkeley, United States
sastry@eecs.berkeley.edu

ABSTRACT

We propose average unfairness as a new measure of fairness in routing games, defined as the ratio between the average latency and the minimum latency experienced by users. This measure is a natural complement to two existing unfairness notions: loaded unfairness, which compares maximum and minimum latencies of routes with positive flow, and user equilibrium (UE) unfairness, which compares maximum latency with the latency of a Nash equilibrium. We show that the worst-case values of all three unfairness measures coincide and are characterized by a steepness parameter intrinsic to the latency function class. We show that average unfairness is always no greater than loaded unfairness, and the two measures are equal only when the flow is fully fair. Besides that, we offer a complete comparison of the three unfairness measures, which, to the best of our knowledge, is the first theoretical analysis in this direction. Finally, we study the constrained system optimum (CSO) problem, where one seeks to minimize total latency subject to an upper bound on unfairness. We prove that, for the same tolerance level, the optimal flow under an average unfairness constraint achieves lower total latency than any flow satisfying a loaded unfairness constraint. We show that such improvement is always strict in parallel-link networks and establish sufficient conditions for general networks. We further illustrate the latter with numerical examples. Our results provide theoretical guarantees and valuable insights for evaluating fairness-efficiency tradeoffs in network routing.

KEYWORDS

Routing games; Fairness; Constrained system optimum

ACM Reference Format:

Pan-Yang Su, Arwa Alanqary, Bryce L. Ferguson, Manxi Wu, Alexandre M. Bayen, and Shankar Sastry. 2026. Average Unfairness in Routing Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/HAKF7947>

*Equal contribution



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/HAKF7947>

1 INTRODUCTION

In the classical routing games model, user behavior naturally leads to equilibria that can be significantly inefficient compared to centrally optimized solutions. While these optimal flows achieve the lowest total latency among the population, such centralized routing can often compromise *fairness*, potentially discriminating against certain users by assigning them disproportionately high latency routes. This fundamental tension between efficiency and fairness has motivated a growing body of work aimed at formalizing measures of unfairness, understanding its trade-offs with efficiency, and developing algorithms for computing flows that minimize the total latency under unfairness constraints.

A key concern in this line of research is defining what it means for a flow to be fair. Much of the literature adopts worst-case fairness measures, which quantify fairness by the experience of the most disadvantaged user. While this perspective provides strong guarantees, it is brittle: a single outlier can dominate the fairness measure, regardless of how many users actually suffer such extreme latency. As a result, worst-case fairness measures may be overly pessimistic and fail to reflect the average user experience.

In this work, we introduce a new *average* measure of unfairness that captures the average envy experienced by users in a given flow. By shifting the focus from the worst-off user to the expected delay in the network, our approach provides a more stable way to reason about fairness in routing games. In light of this newly introduced measure, we revisit the question: *How unfair is optimal routing?* To that end, we prove a tight bound on the average unfairness of an optimal flow for a class of latency functions. We also revisit two canonical notions of worst-case unfairness: the *loaded unfairness*, which is the maximum ratio between two users' experienced latency, and the *user equilibrium unfairness*, which is the maximum ratio between a user's latency at the given flow and the latency they experience at equilibrium. Interestingly, the bound that we derive for the average unfairness matches the bounds for both of these measures, but the lower bound is attained at different instances. Further, for any given instance and a feasible flow exhibiting some level of unfairness, we show that the average unfairness is always strictly less than the corresponding loaded unfairness. This highlights that worst-case measures can overstate the level of unfairness and it motivates the use of average-case measures as user constraints. Besides that, we offer a complete comparison of all the three unfairness measures. To the best of our knowledge, since their introduction by [1], these relationships have not been formally

studied. Our analysis complements prior work [1], which investigated these connections only through numerical experiments, by providing the first rigorous theoretical treatment. Finally, we study the constrained system optimum problem (CSO) under the alternative unfairness measures. For a fixed tolerance, the CSO solution with the average-unfairness constraint can attain a strictly lower cost than the loaded-unfairness solution. We derive verifiable conditions under which this strict improvement is guaranteed. These conditions are expressed in terms of the path support of the loaded-unfairness CSO solution and always hold in parallel-link networks, ensuring strict improvement in that class. On benchmark networks with general topology, our numerical experiments frequently exhibit the same effect.

The main contributions of this paper are summarized as follows.

- (1) We introduce average unfairness as an intuitive and interpretable measure of unfairness and derive its exact bound under general network topologies and latency function classes.
- (2) We provide a comprehensive theoretical comparison between average unfairness and two commonly used measures in the literature (loaded and user equilibrium (UE) unfairness). To the best of our knowledge, this is the first theoretical comparison since the numerical study in [1].
- (3) We analyze the constrained system optimum (CSO) problem under all three unfairness measures. We establish the structure of its solutions in parallel-link networks, extending the result in [12]. Moreover, we show that the CSO solution under average unfairness can strictly outperform that under loaded unfairness, and establish conditions under which this improvement is guaranteed. We further validate these results through simulations.

1.1 Related Work

Unfairness measures and constrained system optimum. The notion of unfairness in routing games and its first formalization were introduced by Jahn et al. [10]. This work proposed a menu of unfairness measures each capturing different aspects of unfairness. One of the most studied among these is the *loaded unfairness*, which measures the maximum ratio between the latencies of used paths. It captures the worst-case envy among users by quantifying how much more latency one user experiences compared to another. This measure has been further analyzed in [3, 6], and a closely related definition was introduced and analyzed in [11]. An analogous measure, the minmax fairness or envy-ratio, was proposed and analyzed for the atomic congestion games model in [4, 9].

Another measure, the *user equilibrium unfairness*, captures the deviation from the Nash equilibrium flow. It is defined as the maximum ratio between the latency of a path used in the current flow and that of a path used at equilibrium. This measure was studied in [15] to bound the unfairness of an optimal flow. Jahn et al. [10] also introduced the *normal unfairness*, which evaluates fairness based on the ratio of the normal length of a given path to that of the shortest path. The normal length refers to any property of the path that is fixed a priori and is independent of the flow. While attractive for its simplicity [1], this measure fails to reflect the actual level of experienced unfairness under realistic routing conditions [2]. An improved notion of normal unfairness that uses the path's

latency at the Nash equilibrium flow as the normal length was proposed and analyzed in [16]. Finally, the *free-flow unfairness*—the least studied of the four—measures the maximum ratio between the latency of a used path and its latency under no congestion. It can be viewed as an intermediary between the normal and loaded unfairness measures.

We note that all these unfairness measures that are analyzed in the context of routing games are worst-case measures, focusing on the extremes of the population. In contrast, the average unfairness measure we propose shifts the focus to quantifying the average envy each user feels toward the most-advantaged user in the population. This approach aligns with some fairness notions studied in the broader resource allocation literature, where population-level measures are common. A notable example is the Gini index [7], which captures average pairwise differences in users' costs normalized by the average cost. A key technical distinction is that the Gini index is based on cost differences, while our measure uses cost ratios. Conceptually, while the Gini index reflects the average envy between all pairs of users, our measure captures the average envy of each user relative to the best-off user, providing a perspective more aligned with the previously studied measures in this literature.

Previous numerical comparisons between these measures [10, 11] highlight the fact that they are not equivalent. In this work, we strengthen the understanding of these relationships by establishing theoretical comparisons between the measures, including the newly introduced average unfairness, whenever they are comparable, and illustrate with examples where such relationships do not hold.

The study of unfairness is motivated by the *constrained system optimum (CSO)* problem which aims at finding flows that minimize the total latency among those flows with bounded unfairness. This problem is of great practical interest as it guides the design of routing algorithms and recommendations. Our average-case unfairness measure gives rise to a new formulation of the CSO model that is more realistic and potentially with better (more efficient) solutions.

Approximate equilibrium and price of anarchy and stability. Unfairness measures are closely related to the concept of approximate equilibria in routing games, though they are motivated by distinct concerns: fairness versus bounded rationality. Despite their different motivations, the definitions of measures and solution sets, as well as the analytical tools used in both lines of research are similar. A key distinction, however, lies in the fact that approximation measures fully characterize the set of exact Nash equilibria as the approximation tolerance goes to zero, a property not shared by any of the unfairness measures commonly studied. This creates subtle differences in the way these measures are defined. For instance, the multiplicative definition of approximate equilibrium [15] compares the cost of a used path to that of any alternative, whereas the loaded unfairness introduced above only compares used paths. Moreover, performance measures developed for approximate equilibria, such as the price of stability (POS) [5], or for tolled equilibria, such as the tolled price of anarchy (POA) [8], offer complementary perspectives useful for the study of fairness-related problems, such as the CSO problem. In particular, the POS, with respect to a suitable unfairness/approximation measure, offer a meaningful upper bound on the values of the corresponding CSO problem for a class of routing games.

2 PRELIMINARIES

We provide the preliminaries of a routing game, following the model in [13, 15].

2.1 The Model

We consider a directed network $G = (V, E)$ with vertex set V , edge set E , and k source-destination vertex pairs $\{s_1, t_1\}, \dots, \{s_k, t_k\}$, where each pair is referred to as a *commodity*. To simplify notation, we let $[k]$ denote the set of indices $\{1, 2, \dots, k\}$. We denote the set of (simple) s_i - t_i paths by \mathcal{P}_i , which is assumed to be non-empty, and define $\mathcal{P} = \cup_{i \in [k]} \mathcal{P}_i$. A *flow* is a function $f : \mathcal{P} \rightarrow \mathbb{R}^+$; for a fixed flow f we define $f_e = \sum_{P \in \mathcal{P}: e \in P} f_P$. We associate a finite and positive *rate* r_i with each pair $\{s_i, t_i\}$, the amount of flow with source s_i and destination t_i ; a flow f is said to be *feasible* if for all i , $\sum_{P \in \mathcal{P}_i} f_P = r_i$. Finally, each edge $e \in E$ is given a load-dependent *latency function* that we denote by ℓ_e . For each $e \in E$, we assume that the latency function ℓ_e is nonnegative, continuous, and nondecreasing. Following [14, Definition 2.5], the following definition will be useful later.

Definition 2.1. A latency function ℓ is standard if $x \cdot \ell(x)$ is convex on $[0, \infty)$.

We will call the triple (G, r, ℓ) an *instance*. The latency of a path P with respect to a flow f is defined as the sum of the latencies of the edges in the path, denoted by $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$ ¹. We define the *cost* $C(f)$ of a flow f in (G, r, ℓ) as the total latency incurred by f ; that is,

$$C(f) = \sum_{P \in \mathcal{P}} \ell_P(f) f_P. \quad (1)$$

By summing over the edges in a path P and reversing the order of summation, we may also write $C(f) = \sum_{e \in E} \ell_e(f_e) f_e$. Also, the total latency of commodity i is $C_i(f) = \sum_{P \in \mathcal{P}_i} \ell_P(f) f_P$.

Finally, given a class \mathcal{L} of latency functions, we define a class of instances $\mathcal{G}(\mathcal{L}) = \{(G, r, \ell) \mid G \in \mathcal{G}_{fin}^{dir}, r > 0, \ell_e \in \mathcal{L}\}$, where \mathcal{G}_{fin}^{dir} is the set of finite directed graphs. A class $\mathcal{G}(\mathcal{L})$ represents all routing games on finite directed graphs with latency functions from the class \mathcal{L} .

2.2 Nash and Optimal Flows

Fix a flow f feasible for an instance (G, r, ℓ) . We say that f is an *optimal flow* if it minimizes $C(f)$. We denote the set of optimal flows by $OPT(G, r, \ell)$. We say that it is a *Nash flow* if for every i and every two s_i - t_i paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $\ell_{P_1}(f) \leq \ell_{P_2}(f)$. With nondecreasing and continuous latency functions, both optimal flows and Nash flows exist².

Given an instance (G, r, ℓ) with standard latency functions, define the *marginal cost function* $\hat{\ell}_e = \frac{d}{dx}(x \cdot \ell_e(x))$. Then, a flow f feasible for (G, r, ℓ) is optimal if and only if it is a Nash flow for $(G, r, \hat{\ell})$; that is, for every i and every two s_i - t_i paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $\hat{\ell}_{P_1}(f) \leq \hat{\ell}_{P_2}(f)$.

¹Note that $\ell_e(f_e)$ only depends on f_e , the amount of traffic on edge e , but $\ell_P(f)$ depends not only on f_P , the amount of traffic on path P , but also on the amount of traffic on each edge $e \in P$. Specifically, $\ell_P(f)$ depends on $\{f_e \mid e \in P\}$.

²Continuity alone suffices for the existence of an optimal flow due to the extreme value theorem, but the existence of a Nash flow needs both assumptions; see [15, Lemma 2.6].

2.3 Unfairness

We revisit two load-dependent unfairness measures introduced in [10] and define the average unfairness measure. Note that we define $0/0 := 1$, so a flow where all traffic incurs zero latency has an unfairness of 1 (totally fair).

Definition 2.2 (Loaded unfairness). Given an instance (G, r, ℓ) and a feasible flow f , the *loaded unfairness* $U^L(G, r, \ell, f)$ is defined as

$$U^L(G, r, \ell, f) = \max \left\{ \frac{\ell_P(f)}{\ell_Q(f)} \mid P, Q \in \mathcal{P}_i, f_P, f_Q > 0, i \in [k] \right\}. \quad (2)$$

We denote $U_i^L(G, r, \ell, f) = \max \left\{ \frac{\ell_P(f)}{\ell_Q(f)} \mid P, Q \in \mathcal{P}_i, f_P, f_Q > 0 \right\}$ for commodity i 's loaded unfairness. For a given instance, we denote its loaded unfairness $U^L(G, r, \ell) := \sup_{f^* \in OPT(G, r, \ell)} U^L(G, r, \ell, f^*)$. Given a class of latency functions \mathcal{L} , we denote its loaded fairness $U^L(\mathcal{L}) = \sup_{(G, r, \ell) \in \mathcal{G}(\mathcal{L})} U^L(G, r, \ell)$.

REMARK 1. In defining $U^L(G, r, \ell)$, we take the supremum rather than the infimum to capture the worst-case scenario, although, in practice, one might seek an optimal flow that minimizes unfairness.

REMARK 2. Different flows, particularly optimal flows $f^* \neq \tilde{f}^*$, that induce the same edge flows $f_e^* = \tilde{f}_e^*, \forall e \in E$, can have different values of loaded unfairness $U^L(f^*) \neq U^L(\tilde{f}^*)$. To illustrate this, consider the network in Fig. 1 which consists of two Pigou networks connected in series. When routing one unit of traffic, any flow that induces traffic of $1/2$ to each edge is optimal, but splitting the unit demand equally between paths 1 and 2 yields unfairness 2 (worst case unfairness for instances with linear latency functions [6]), while splitting between paths 3 and 4 yields unfairness 1 (totally fair). This example shows that even though these flows have the same cost, they induce different levels of unfairness.




Figure 1: Two Pigou networks connected in series.

Definition 2.3 (UE unfairness). Given an instance (G, r, ℓ) and a feasible flow f , the *user equilibrium (UE) unfairness* $U^{UE}(G, r, \ell, f)$ is defined as

$$U^{UE}(G, r, \ell, f) = \max \left\{ \frac{\ell_P(f)}{\ell_Q(f^{NE})} \mid P, Q \in \mathcal{P}_i, f_P, f_Q^{NE} > 0, i \in [k] \right\}, \quad (3)$$

where f^{NE} is a Nash flow.

We denote $U_i^{UE}(G, r, \ell, f) = \max\{\frac{\ell_P(f)}{\ell_Q(f^{NE})} | P, Q \in \mathcal{P}_i, f_P, f_Q^{NE} > 0\}$ for commodity i 's UE unfairness. For a given instance, we denote its UE unfairness $U^{UE}(G, r, \ell) := \sup_{f^* \in OPT(G, r, \ell)} U^{UE}(G, r, \ell, f^*)$. Given a class of latency functions \mathcal{L} , we denote its UE fairness $U^{UE}(\mathcal{L}) = \sup_{(G, r, \ell) \in \mathcal{G}(\mathcal{L})} U^L(G, r, \ell)$.

REMARK 3. Recall from [15, Lemma 2.6] that any instance (G, r, ℓ) with continuous and nondecreasing latency functions admits a Nash flow. Also, if f^{NE} and \hat{f}^{NE} are Nash flows, then $C_i(f^{NE}) = C_i(\hat{f}^{NE})$ for any commodity i . This fact, together with the property that all paths with positive flow in a Nash flow have equal latency, ensures that $\ell_Q(f^{NE})$ in (3) is well-defined.

Now, we introduce the average unfairness measure, motivated by two equivalent interpretations: expected unfairness and average envy. Consider an instance (G, r, ℓ) and a feasible flow f . For simplicity, consider a single commodity. A commuter traveling on a path $P \in \mathcal{P}$ experiences ex post unfairness $\frac{\ell_P(f)}{\ell_Q(f)}$, where Q is the minimum-latency path with a positive flow. A natural definition of the envy of a commuter on a path P is precisely this ex post unfairness ratio $\frac{\ell_P(f)}{\ell_Q(f)}$. Thus, loaded unfairness captures the maximum envy (see (2)), while average unfairness captures the average envy across travelers, given by

$$\frac{C(f)}{r \ell_Q(f)} = \sum_{P \in \mathcal{P}} \frac{f_P}{r} \cdot \frac{\ell_P(f)}{\ell_Q(f)}.$$

Beyond this average-envy interpretation, the ratio $C(f)/(r \ell_Q(f))$ can also be viewed as the expected unfairness experienced by a commuter when f is interpreted as an ex ante route assignment. These observations motivate the following definition of the average unfairness measure.

Definition 2.4 (Average unfairness). Given an instance (G, r, ℓ) and a feasible flow f , the average unfairness $U^A(G, r, \ell, f)$ is defined as

$$U^A(G, r, \ell, f) = \max \left\{ \frac{C_i(f)}{r_i \ell_Q(f)} | Q \in \mathcal{P}_i, f_Q > 0, i \in [k] \right\}. \quad (4)$$

We denote $U_i^A(G, r, \ell, f) = \max\{\frac{C_i(f)}{r_i \ell_Q(f)} | Q \in \mathcal{P}_i, f_Q > 0\}$ for commodity i 's average unfairness. For a given instance, we denote its average unfairness $U^A(G, r, \ell) := \sup_{f^* \in OPT(G, r, \ell)} U^A(G, r, \ell, f^*)$. Given a class of latency functions \mathcal{L} , we denote its average fairness $U^A(\mathcal{L}) = \sup_{(G, r, \ell) \in \mathcal{G}(\mathcal{L})} U^A(G, r, \ell)$.

REMARK 4. We view different commodities as incomparable and take the maximum across different commodities in (4). One can also define the average unfairness to be the average of different commodities' average unfairness. All the results in this work hold for both definitions.

3 UNFAIRNESS BOUNDS AND COMPARISONS

3.1 Unfairness Bounds

In this section, we will prove the main result that gives the exact bound of average unfairness of a class of instances under general conditions. The bound is reminiscent of the results in [6, Theorem 4.2] and [13, Theorem 3.1], which we unify here for completeness.

As with those results, the bound depends on the steepness of the class of latency functions.

Definition 3.1 (Steepness). Let \mathcal{L} denote a class of differentiable latency functions. For $\ell \in \mathcal{L}$, define the steepness $\gamma(\ell)$ as $\gamma(\ell) = \sup_{x>0} \frac{\hat{\ell}(x)}{\ell(x)}$. Define $\gamma(\mathcal{L}) = \sup_{\ell \in \mathcal{L}} \gamma(\ell)$.

We also consider the following set of assumptions.

ASSUMPTION 1. We assume that \mathcal{L} is a family of differentiable and standard latency functions.

ASSUMPTION 2. We assume that \mathcal{L} includes all constant functions.

THEOREM 3.2. Let \mathcal{L} be a class of latency functions satisfying Assumption 1 and Assumption 2. Then $U^L(\mathcal{L}) = U^{UE}(\mathcal{L}) = U^A(\mathcal{L}) = \gamma(\mathcal{L})$.

PROOF. The equalities $U^L(\mathcal{L}) = \gamma(\mathcal{L})$ and $U^{UE}(\mathcal{L}) = \gamma(\mathcal{L})$ are proved in [6, Theorem 4.2] and [13, Theorem 3.1], respectively. The equality for the average unfairness follows from Lemma 3.3 and Lemma 3.5 proved below. \square

REMARK 5. Assumption 1 alone suffices for the three unfairness measures to be upper bounded by $\gamma(\mathcal{L})$, but both assumptions are required for the lower bound.

REMARK 6. We highlight two points about the results in [6, Theorem 4.2] and [13, Theorem 3.1]. First, although the assumption of standard latency functions is not stated explicitly in [6, Theorem 4.2], it cannot be dispensed with. Second, while [13, Theorem 3.1] only considers the single-commodity case, the result extends to multiple commodities.

LEMMA 3.3. Let \mathcal{L} be a family of latency functions satisfying Assumption 1. Then $U^A(\mathcal{L}) \leq \gamma(\mathcal{L})$.

PROOF. Fix any instance (G, r, ℓ) and an optimal flow f^* . Consider any $i \in [k]$. Since f^* is optimal, we get $\hat{\ell}_P(f^*) \leq \hat{\ell}_Q(f^*)$ for any $P, Q \in \mathcal{P}_i$ with $f_P^* > 0$. Let Q be a minimum-latency path with a positive flow for commodity i , multiply both sides by f_P^* , and sum over all paths:

$$\sum_{P \in \mathcal{P}_i} f_P^* \hat{\ell}_P(f^*) \leq \left(\sum_{P \in \mathcal{P}_i} f_P^* \right) \cdot \hat{\ell}_Q(f^*) = r_i \hat{\ell}_Q(f^*). \quad (5)$$


Combining the results, we get

$$C_i(f^*) \leq \text{LHS of (5)} \leq \text{RHS of (5)} \leq \gamma(\mathcal{L}) \cdot r_i \ell_Q(f^*), \quad (6)$$

where we use $\hat{\ell}_e(f_e^*) = \ell_e(f_e^*) + f_e^* \ell'_e(f_e^*) \geq \ell_e(f_e^*)$ in the first inequality and $\hat{\ell}_Q(f^*) = \sum_{e \in Q} \hat{\ell}_e(f_e^*) \leq \gamma(\mathcal{L}) \sum_{e \in Q} \ell_e(f_e^*) = \gamma(\mathcal{L}) \ell_Q(f^*)$ in the third inequality. Since (6) holds for any commodity i , $U^A(\mathcal{L}) \leq \gamma(\mathcal{L})$. \square

REMARK 7. Another way to prove Lemma 3.3 is to invoke the inequality $U^L(G, r, \ell, f) \leq \gamma(\mathcal{L})$ in [6, Theorem 4.2] and the inequality $U^A(G, r, \ell, f) \leq U^L(G, r, \ell, f)$ in Proposition 3.7, which will be proved later without using Lemma 3.3.

In proving the lower bound, we make use of the Pigou network in Fig. 2. Before presenting the general lower bound in Lemma 3.5, we provide a motivating example of the average unfairness of the optimal flow when latency functions are polynomials.


Figure 2: Pigou network.

Example 3.4. Construct (G, r, ℓ) as follows: consider one unit of traffic ($r = 1$) traveling through the Pigou network in Fig. 2, with $\ell(x) = x^n$ and $c = \epsilon$ where $\epsilon, n > 0$. Solving $(x^{n+1} + (1-x)\epsilon)' = 0$, we get $x = (\frac{\epsilon}{n+1})^{\frac{1}{n}}$, so

$$U^A(G, r, \ell) \geq \frac{\frac{\epsilon}{n+1} (\frac{\epsilon}{n+1})^{\frac{1}{n}} + \epsilon(1 - (\frac{\epsilon}{n+1})^{\frac{1}{n}})}{\frac{\epsilon}{n+1}} \xrightarrow{\epsilon \rightarrow 0} n + 1. \quad (7)$$

LEMMA 3.5. *Let \mathcal{L} be a family of latency functions satisfying Assumption 1 and Assumption 2. Then $U^A(\mathcal{L}) \geq \gamma(\mathcal{L})$.*

PROOF. Consider $k = 1$. For any $\delta \in (0, \gamma(\mathcal{L}))$, there are $\ell \in \mathcal{L}$ and $\epsilon > 0$ such that $\hat{\ell}(\epsilon) \geq (\gamma(\mathcal{L}) - \delta)\ell(\epsilon)$.

Consider r units of traffic traveling through the Pigou network with latency functions $\hat{\ell}(\epsilon)$ and $\ell(x)$, respectively, and $r > \epsilon$. Solving $(x\ell(x) + \hat{\ell}(\epsilon)(r-x))' = 0$, we get $\hat{\ell}(x) = \hat{\ell}(\epsilon)$, so we assume $x = \epsilon$ below.

$$U^A(\mathcal{L}) \geq \lim_{r \rightarrow \infty} \frac{\epsilon\ell(\epsilon) + \hat{\ell}(\epsilon)(r-\epsilon)}{r\ell(\epsilon)} = \lim_{r \rightarrow \infty} \frac{\epsilon}{r} + \frac{\hat{\ell}(\epsilon)(r-\epsilon)}{r\ell(\epsilon)}. \quad (8)$$

The first term goes to 0, so we focus on the second term.

$$\frac{\hat{\ell}(\epsilon)(r-\epsilon)}{r\ell(\epsilon)} \geq \frac{\hat{\ell}(\epsilon)(r-\epsilon)}{r\frac{\hat{\ell}(\epsilon)}{\gamma(\mathcal{L})-\delta}} = (1 - \frac{\epsilon}{r})(\gamma(\mathcal{L}) - \delta) \xrightarrow{r \rightarrow \infty} \gamma(\mathcal{L}) - \delta. \quad (9)$$

Letting $\delta \rightarrow 0$ finishes the proof. \square

COROLLARY 3.6. *If \mathcal{L} is the set of polynomials with degree at most n and nonnegative coefficients, $U^A(\mathcal{L}) = n + 1$.*

PROOF. Since $(x \cdot x^n)' = (n+1)x^n$, the result follows from Theorem 3.2. \square

3.2 Comparisons

It is informative to compare the values that the three unfairness measures take at the same flow. For the loaded and average unfairness measures, the following proposition characterizes this comparison.

PROPOSITION 3.7. *Given any instance (G, r, ℓ) and feasible flow f , one of the following holds.*

- (1) $U^L(G, r, \ell, f) > U^A(G, r, \ell, f) > 1$.
- (2) $U^L(G, r, \ell, f) = U^A(G, r, \ell, f) = 1$.


PROOF. Fix an instance (G, r, ℓ) and drop the dependence. For any feasible flow f and commodity i , since $r_i \max_{P \in \mathcal{P}_i, f_P > 0} \ell_P(f) \geq \sum_{P \in \mathcal{P}_i} f_P \ell_P(f)$, we have $U_i^L \geq U_i^A$, so $U^L \geq U^A$. Below, we show

$$\begin{aligned} U^L = U^A &\iff U^A = 1. \\ U^L = U^A &\iff \exists i^* \in [k] \text{ s.t. } U_{i^*}^L = U_{i^*}^A \geq U^A \\ &\iff r_{i^*} \max_{P \in \mathcal{P}_{i^*}, f_P > 0} \ell_P(f) = \sum_{P \in \mathcal{P}_{i^*}} f_P \ell_P(f) \\ &\iff \max_{P \in \mathcal{P}_{i^*}, f_P > 0} \ell_P(f) = \min_{P \in \mathcal{P}_{i^*}, f_P > 0} \ell_P(f) \\ &\iff U_{i^*}^A = U^A = 1. \end{aligned} \quad (10)$$

\square

A result similar to Proposition 3.7 cannot be established between the UE unfairness and the loaded unfairness as there is no general relationship between the two measures (see Example 3.8 and Example 3.10). It is also possible that at a given flow, one of the measures equals 1 while the other does not (see Example 3.8 and Example 3.9). Moreover, as noted in [10, Section 2.2], the loaded unfairness is always greater than or equal to one, whereas UE unfairness can take any nonnegative value (see Example 3.9).

Example 3.8. Construct (G, r, ℓ) as follows: consider one unit of traffic ($r = 1$) traveling through the network in Fig. 1. Consider the (optimal) flow f^* that routes half of the traffic on the upper path (path 1) and the other half on the lower path (path 2). Then $U^L(G, r, \ell, f^*) = 2 > U^{UE}(G, r, \ell, f^*) = 1$.


Figure 3: Braess's network.

Example 3.9. Construct (G, r, ℓ) as follows: consider one unit of traffic ($r = 1$) traveling through the Braess's network in Fig. 3. Consider the (optimal) flow f^* that routes half of the traffic on the upper path (path 1) and the other half on the lower path (path 3). Then $U^L(G, r, \ell, f^*) = 1 > U^{UE}(G, r, \ell, f^*) = \frac{3}{4}$.

Example 3.10. Construct (G, r, ℓ) as follows: consider the multi-commodity network in Fig. 4. Consider the (optimal) flow f^* that routes all the traffic of the first commodity along the path $s_1 - t$ and all the traffic of the second commodity along the only path $s_2 - t$. Note that the Nash flow routes all the traffic of the first commodity through the path $s_1 - s_2 - t$. Thus, $U^{UE}(G, r, \ell, f^*) = 3/2 > U^L(G, r, \ell, f^*) = 1$.

However, if we restrict our attention to optimal flows in single-commodity instances, we obtain a relationship similar to that in Proposition 3.7. Although both unfairness measures have been extensively studied in the literature for the past decade, to the best of our knowledge, this kind of relationship has never been theoretically established.

PROPOSITION 3.11. *Given any single-commodity ($k = 1$) instance (G, r, ℓ) and optimal flow f^* , one of the following holds.*




Figure 4: Multi-commodity network.

- (1) $U^L(G, r, \ell, f^*) > U^{UE}(G, r, \ell, f^*)$.
- (2) $U^L(G, r, \ell, f^*) = U^{UE}(G, r, \ell, f^*) = 1$.

PROOF. Fix an instance (G, r, ℓ) and drop the dependence. For any optimal flow f^* and Nash flow f^{NE} , we have the relationship $\min_{P \in \mathcal{P}_i, f_P^* > 0} \ell_P(f^*) \leq \max_{P \in \mathcal{P}_i, f_P^{NE} > 0} \ell_P(f^{NE})$; for otherwise, f^{NE} gives a lower total cost, a contradiction. Thus, $U^L \geq U^{UE}$.

Below, we show $U^L = U^{UE} \iff (U^{UE} = 1 \text{ and } U^L = 1)$. We only need to show the only if direction.

$$\begin{aligned}
 U^L = U^{UE} &\iff \min_{P \in \mathcal{P}_i, f_P^* > 0} \ell_P(f^*) = \max_{P \in \mathcal{P}_i, f_P^{NE} > 0} \ell_P(f^{NE}) \\
 &\Rightarrow \max_{P \in \mathcal{P}_i, f_P^* > 0} \ell_P(f^*) = \min_{P \in \mathcal{P}_i, f_P^* > 0} \ell_P(f^*) \\
 &= \max_{P \in \mathcal{P}_i, f_P^{NE} > 0} \ell_P(f^{NE}) \\
 &\Rightarrow U^{UE} = 1 \text{ and } U^L = 1,
 \end{aligned} \tag{11}$$

where the second implication holds by contradiction; if it does not hold, f^* cannot be optimal. \square

REMARK 8. Proposition 3.11 does not hold for multiple commodities (see Example 3.10). However, it holds for another version of the definitions when we are allowed to compare the latencies of different commodities. Formally speaking, if we replace the condition $P, Q \in \mathcal{P}_i$ in (2) and (3) with $P, Q \in \mathcal{P}$, Proposition 3.11 holds for the modified definitions with multiple commodities.

The comparisons also give rise to the following result, exemplified by Example 3.4 and the example in [13, Section 1], where the upper bound for average unfairness and UE unfairness are approached, but not attained, when the respectively defined problem parameters ϵ goes to zero.

COROLLARY 3.12. Let \mathcal{L} be a family of latency functions.

- (1) If $U^A(\mathcal{L}) > 1$, then $U^A(G, r, \ell, f^*) < U^A(\mathcal{L})$ for any $(G, r, \ell) \in \mathcal{G}(\mathcal{L})$ and $f^* \in OPT(G, r, \ell)$.
- (2) If $k = 1$ and $U^{UE}(\mathcal{L}) > 1$, then $U^{UE}(G, r, \ell, f^*) < U^{UE}(\mathcal{L})$ for any $(G, r, \ell) \in \mathcal{G}(\mathcal{L})$ and $f^* \in OPT(G, r, \ell)$.

PROOF. We prove the first part. If an instance (G, r, ℓ) and optimal flow f^* attain the upper bound $U^A(\mathcal{L})$, then we have

$$U^L(G, r, \ell, f^*) = U^A(G, r, \ell, f^*) > 1, \tag{12}$$

violating Proposition 3.7. The second part follows similarly from Proposition 3.11. \square

Finally, we show that there is no general relationship between the UE unfairness and average unfairness; specifically, they can be drastically different in different instances.

PROPOSITION 3.13. Given any instance (G, r, ℓ) and feasible flow f , there is no relationship between $U^{UE}(G, r, \ell, f)$ and $U^A(G, r, \ell, f)$.

PROOF. Consider Example 3.4, where $\lim_{\epsilon \rightarrow 0} U^A(G, r, \ell, f^*) = n+1$ with the optimal flow f^* . However, we have $U^{UE}(G, r, \ell, f^*) = 1$ for every $\epsilon > 0$. In contrast, consider the Pigou network in Fig. 2 with $\ell(x) = x^n$ and $c = (n+1)(1-\epsilon)$ (see [13]). Then $U^A(G, r, \ell, f^*) = (n+1) - n(1-\epsilon)^{\frac{1}{n}} \xrightarrow{\epsilon \rightarrow 0} 1$, but $U^{UE}(G, r, \ell, f^*) = \frac{(n+1)(1-\epsilon)}{1} \xrightarrow{\epsilon \rightarrow 0} n+1$. \square

4 CONSTRAINED SYSTEM OPTIMUM (CSO)

In Section 3.2, we analyzed the worst-case unfairness for a class of instances across the different measures and compared the values of unfairness produced by these measures at a given instance and flow. A closely related question is that of the constrained system optimum problem first proposed in [10]. In this model, the goal is to find a feasible flow f that minimizes the total latency with constraints on the unfairness it induces among the population. More formally, the CSO problem can be formulated as

$$(CSO) \quad \arg \min_{f: U(f) \leq 1+\beta} C(f), \tag{13}$$

where U is the unfairness measure of choice and $\beta \geq 0$ is the unfairness tolerance. Different unfairness measures give different solutions to this problem. In view of Proposition 3.7, in this section, we focus on comparing the solution sets of the CSO problem under loaded unfairness and average unfairness constraints. We denote by $OPT^L(G, r, \ell, \beta)$ and $OPT^A(G, r, \ell, \beta)$ the sets of optimal flows under the loaded unfairness and average unfairness constraints, respectively. Also, denote the cost of the respective sets by

$$\begin{aligned}
 C^L(G, r, \ell, \beta) &= \min_{f: U^L(f) \leq 1+\beta} C(f), \\
 C^A(G, r, \ell, \beta) &= \min_{f: U^A(f) \leq 1+\beta} C(f).
 \end{aligned} \tag{14}$$

An immediate consequence of Proposition 3.7 is Corollary 4.1.

COROLLARY 4.1. Given any instance (G, r, ℓ) and unfairness tolerance $\beta \geq 0$, $C^A(G, r, \ell, \beta) \leq C^L(G, r, \ell, \beta)$.

In some instances, the average unfairness constraint can strictly improve upon loaded unfairness; that is, Corollary 4.1 holds with strict inequality (see Example 4.2). However, this is not always the case (see Example 4.3).

Example 4.2. Construct (G, r, ℓ) as follows: consider one unit of traffic ($r = 1$) traveling through the Pigou network in Fig. 2 with $\ell(x) = x$ and $c = 1 + \beta$, where $\beta \in (0, 1)$. The optimal flow of this instance f^* routes $\frac{1+\beta}{2}$ on the upper edge and $\frac{1-\beta}{2}$ on the lower edge with $C(f^*) = \frac{3+2\beta-\beta^2}{4} < 1$. The optimal flow under the loaded unfairness $\{f^L\} = OPT^L(G, r, \ell, \beta)$ routes all traffic on the upper edge with unfairness of $U^L(f^L) = 1$ and cost $C^L(G, r, \ell, \beta) = 1$. However, for the average unfairness, we can improve the cost of f^L by moving a small amount of traffic to the lower edge without violating the unfairness constraint.

Example 4.3. Construct (G, r, ℓ) as follows: consider $3/4$ unit of traffic ($r = 3/4$) traveling through the Braess's network in Fig. 3.

The optimal flow f^* distributes the traffic evenly across all three paths ($f_{P_1}^* = f_{P_2}^* = f_{P_3}^* = \frac{1}{4}$) with a total cost of 1. Consider a small unfairness tolerance $\beta > 0$. Simple calculation shows that the flow f^β defined by $f_{P_1}^\beta = f_{P_3}^\beta = 3/8$ and $f_{P_2}^\beta = 0$ is the unique optimal flow under both unfairness measures, that is $OPT^L(G, r, \ell, \beta) = OPT^A(G, r, \ell, \beta) = \{f^\beta\}$. Therefore, while f^β is not optimal in terms of total latency, and its loaded unfairness $U^L(G, r, \ell, f^\beta) = 1$ is strictly less than $1 + \beta$, there is no feasible direction to improve upon it, even when considering average unfairness.

4.1 Conditions for Strict Improvement

We derive a sufficient condition for the constrained system optimal flow under average unfairness to achieve a strictly lower cost than that under loaded unfairness; that is, $C^A(G, r, \ell, \beta) < C^L(G, r, \ell, \beta)$. This strict difference in cost is notable for two reasons: i) Proposition 3.7 compares the loaded and average unfairness of the unconstrained optimal flow while here we compare the cost of two different flows with the same value of unfairness in each respective measure, thus providing greater nuance to the comparison of the two measures, and ii) should practitioners seek to solve the CSO problem in an effort to identify desirable flows, if the loaded unfairness measure is used, it can provide flows with greater envy across the population while also having higher cost. The condition for the strict difference in cost can be checked by looking at the paths utilized under the solution $f \in OPT^L(G, r, \ell, \beta)$.

We first give some intuition. Consider a flow $f \in OPT^L(G, r, \ell, \beta)$ that is not a system optimal flow (i.e., $f \notin OPT(G, r, \ell)$). Then there exists an improvement direction, that is, a feasible flow g such that $C(\delta g + (1 - \delta)f) < C(f) \forall \delta \in (0, 1]$. This follows directly from the convexity of the (unconstrained) system optimum problem (under Assumption 1). One only needs to check if there exists $\delta \in (0, 1]$ such that $h = \delta g + (1 - \delta)f$ does not violate the average unfairness constraint $U^A(h) \leq 1 + \beta$. This is equivalent to checking the continuity of U^A around f , along with the fact that $U^A(f) < 1 + \beta$ (see Proposition 3.7).

We show that a sufficient condition for this to hold is if f utilizes any of the minimum-latency paths for each commodity. To formalize this, we introduce the following notation for the set of utilized paths $P^+(f)$, the set of minimum-latency paths per commodity $P_{\min}^i(f)$, and the set of minimum-latency *utilized* paths per commodity $P_{\min}^{i,+}(f)$:

$$P^+(f) := \{P \in \mathcal{P} \mid f_P > 0\}. \quad (15)$$

$$P_{\min}^i(f) := \arg \min_{P \in \mathcal{P}_i} \ell_P(f). \quad (16)$$

$$P_{\min}^{i,+}(f) := \arg \min_{P \in \mathcal{P}_i, f_P > 0} \ell_P(f). \quad (17)$$

PROPOSITION 4.4. *Given any instance (G, r, ℓ) satisfying Assumption 1 and $\beta > 0$, if there exists $f \in OPT^L(G, r, \ell, \beta)$ such that $P^+(f) \cap P_{\min}^i(f) \neq \emptyset$ for all $i \in [k]$, then one of the following holds*

- (1) $C^L(G, r, \ell, \beta) = C^A(G, r, \ell, \beta) = C^{OPT}(G, r, \ell)$.
- (2) $C^L(G, r, \ell, \beta) > C^A(G, r, \ell, \beta) \geq C^{OPT}(G, r, \ell)$.

PROOF. We state the proof for the single-commodity case, which extends straightforwardly to instances with multiple commodities. If $C^L(G, r, \ell, \beta) > C^{OPT}(G, r, \ell)$, then there exists a feasible flow g

such that $C(\delta g + (1 - \delta)f) < C(f), \forall \delta \in (0, 1]$. By Proposition 3.7, we know that $U^A(f) < 1 + \beta$. We show that $P^+(f) \cap P_{\min}(f) \neq \emptyset$ implies the continuity of the average unfairness measure U^A around f , which then implies existence of δ small enough such that $U^A(\delta g + (1 - \delta)f) \leq 1 + \beta$.

By our assumption that $P^+(f) \cap P_{\min}(f) \neq \emptyset$, it follows that $\min_{P \in P^+(f)} \ell_P(f) = \min_{P \in \mathcal{P}} \ell_P(f)$, and we denote this common value by c^* . Let $h = \delta g + (1 - \delta)f$. By continuity of the latency functions, the following hold when δ is sufficiently small.

- (1) $\ell_P(h) < \ell_Q(h), \forall P \in P_{\min}(f), Q \notin P_{\min}(f)$.
- (2) $P^+(h) \cap P_{\min}(f) \neq \emptyset$ (since $P^+(f) \cap P_{\min}(f) \neq \emptyset$).

Thus, $P_{\min}^+(h) \subseteq P_{\min}(f)$, and we can make $\min_{P \in P^+(h)} \ell_P(h)$ arbitrarily close to c^* . \square

The condition in Proposition 4.4 need not hold in general networks: a CSO solution may not utilize a minimum-latency path like in Example 4.3. It is, however, guaranteed to hold for parallel-link networks (i.e., single-commodity graphs with two nodes connected by a finite set of parallel edges) as we show in Corollary 4.5.

COROLLARY 4.5. *Given any parallel-link instance (G, r, ℓ) satisfying Assumption 1 and given any $\beta > 0$, one of the following holds.*

- (1) $C^L(G, r, \ell, \beta) = C^A(G, r, \ell, \beta) = C^{OPT}(G, r, \ell)$.
- (2) $C^L(G, r, \ell, \beta) > C^A(G, r, \ell, \beta) \geq C^{OPT}(G, r, \ell)$.

PROOF. It suffices to show that any flow $f \in OPT^L(G, r, \ell, \beta)$ utilizes the minimum-latency path: $P^+(f) \cap P_{\min}(f) \neq \emptyset$. The rest follows from Proposition 4.4.

Assume that there are n edges with $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$, and the total flow is r . Suppose that a flow $f \in OPT^L(G, r, \ell, \beta)$ does not utilize any minimum-latency path; that is, $P^+(f) \cap P_{\min}(f) = \emptyset$, and let P_1 be one such unutilized minimum-latency path.

We then construct a new flow g by iteratively transferring flows from the maximum-latency paths to P_1 until one of the following conditions is met. For simplicity of notation, we abbreviate a path P_j as j in what follows.

- (1) $\ell_1(g_1) = \arg \min_{j \in \{2, 3, \dots, n\}, g_j > 0} \ell_j(g_j)$.
- (2) $g_1 = r$, meaning that all the flows are moved to P_1 .

Then, the transformed flow g satisfies the loaded unfairness constraint since either all utilized paths have the same latency or the minimum latency is unchanged while the maximum latency decreases. Also, the following relationships hold.

- (1) $\ell_1(g_1) \leq \arg \min_{j \in \{2, 3, \dots, n\}, f_j > 0} \ell_j(f_j)$.
- (2) $\ell_i(g_i) \leq \ell_i(f_i), \forall i \in \{2, 3, \dots, n\}$.

Let $y = f - g$ be the difference of the two flows. We will show that $C(f) > C(g)$, a contradiction to f being optimal.

$$\begin{aligned} C(f) - C(g) &= \sum_{i=2}^n f_i \ell_i(f_i) - \left(\sum_{i=2}^n g_i \ell_i(g_i) + \left(\sum_{i=2}^n (y_i) \right) \cdot \ell_1 \left(\sum_{j=2}^n (y_j) \right) \right) \\ &= \left(\sum_{i=2}^n g_i \cdot (\ell_i(f_i) - \ell_i(g_i)) \right) + \left(\sum_{i=2}^n y_i \cdot (\ell_i(f_i) - \ell_1 \left(\sum_{j=2}^n (y_j) \right)) \right) \\ &> 0, \end{aligned} \quad (18)$$

where the first term is nonnegative since the latency functions are nondecreasing, and the second term is positive by our construction of the transformed flow g . \square

The construction in Corollary 4.5 is in fact more general. In particular, the same argument can be used to show the following stronger condition for parallel-link networks:

$$\max_{P \in \mathcal{P}, f_P > 0} \ell_P(f_P) \leq \min_{P \in \mathcal{P}, f_P = 0} \ell_P(0) \quad (19)$$

for any solution f of the CSO problem under any of the three unfairness measures. This has an important implication on the set of used paths under any CSO flow as stated in Corollary 4.6³.

COROLLARY 4.6. *Consider a parallel-link instance (G, r, ℓ) with n edges. For any $\beta \geq 0$, choice of unfairness measure $U \in \{A, L, UE\}$ and solution $f \in OPT^U(G, r, \ell, \beta)$, there exists an ordering of the links such that $\ell_1(0) \leq \ell_2(0) \leq \dots \leq \ell_n(0)$ for which the set of utilized paths has the form $P^+(f) = [m]$ for some $m \leq n$.*

Corollary 4.6 extends the characterization in [12] for the set of utilized paths by the Nash and optimal flows to the CSO flows under all three unfairness measures and general latency functions (that are nonnegative, continuous, and nondecreasing). We note here that when multiple paths have identical zero-flow latencies, different optimal flows may require different orderings of the paths.

4.2 Numerical Experiments

We conduct numerical experiments on general networks to compare the CSO cost under loaded, average, and UE unfairness. While strict improvement in the CSO cost under average unfairness, compared to the loaded unfairness, is not guaranteed, our experiments show that it occurs across many benchmark networks. Further, the experiments show that the solution under average unfairness does not uniformly dominate the UE-unfairness solution. This is consistent with Proposition 3.13, which establishes the lack of general ordering between the two measures in general networks.

To obtain approximate CSO solutions, we use the algorithm proposed by [11], which solves a sequence of convex programs whose objective interpolates between the potential function of the routing game and the total cost function:

$$\alpha C(f) + (1 - \alpha) \sum_{e \in E} \int_0^{f_e} \ell_e(s) ds. \quad (20)$$

The interpolation parameter $\alpha \in [0, 1]$ controls the tradeoff between fairness and efficiency: $\alpha = 0$ yields the Nash flow, while $\alpha = 1$ yields the system optimum.

We vary α in increments of 0.01 and compute a flow f^α for each value; this procedure is agnostic to the unfairness measure. For each resulting flow f^α , we compute all unfairness values. This allows us to trace the Pareto frontier between fairness and total cost for the three measures. For clarity of presentation, instead of absolute total cost values, we use the inefficiency ratio $\rho(f) = C(f^\alpha)/C(f^*)$ where $f^* \in OPT(G, r, \ell)$. The implementation details follow [11], except that our unfairness measures depend explicitly on path assignments (see Remark 2).

The experiments are conducted on four benchmark networks [18] described in Table 1. We use the Bureau of Public Roads (BPR) link latency functions [17] $\ell_e(f_e) = \xi_e \left(1 + a \left(\frac{f_e}{\kappa_e}\right)^4\right)$, where ξ_e, κ_e

³The construction of g does not require the CSO problem to be convex, so Assumption 1 is not needed for Corollary 4.6.

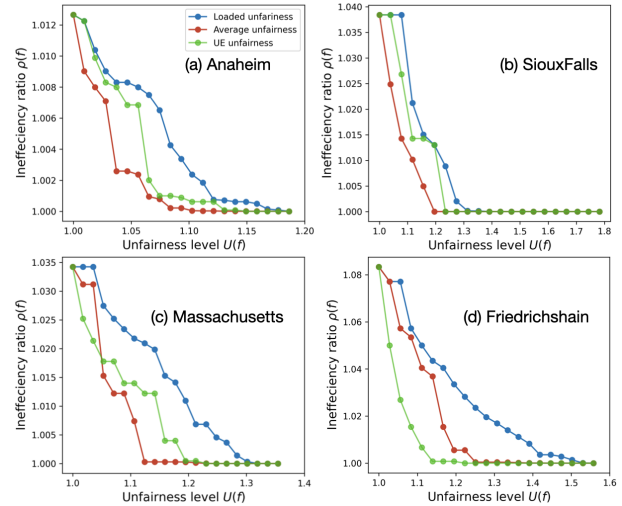


Figure 5: Comparison of CSO costs across varying unfairness levels under loaded (blue), average (red), and UE (green) unfairness measures for four benchmark networks.

are properties of the link and $a = 0.15$. The results are summarized in Fig. 5. The results show that the solutions obtained under the average and UE unfairness do not admit a consistent dominance ordering across networks and unfairness levels. In network (d), the UE unfairness yields more efficient solutions at all unfairness levels, whereas in networks (a) and (b), the average unfairness yields more efficient solutions. In contrast, across all tested networks and unfairness levels, average unfairness consistently achieves a strictly lower cost than loaded unfairness, except for a single data point in network (d). These experiments illustrate the benefit of using the average unfairness measure as the user constraint in the CSO problem in real-world networks.

Table 1: Attributes of the used benchmark networks.

Network Name	attributes		
	$ V $	$ E $	k
Anaheim	416	914	1406
Sioux Falls	24	76	528
Massachusetts	74	258	1113
Friedrichshain	224	523	506

5 CONCLUSION AND FUTURE WORK

We introduced average unfairness as a new and intuitive measure in routing games. We derived a tight worst-case bound on the unfairness of optimal flows under this new measure, and compared it to commonly studied unfairness notions. Our analysis has implications for the constrained system optimum (CSO) problem, showing that average unfairness can lead to more efficient solutions than loaded unfairness constraints. An interesting direction for future work is to study bounds on the value of the CSO problem under average unfairness constraints, in the spirit of price of stability (POS) analyses (cf. [5]). Finally, we aim to explore broader applications of average unfairness in other contexts, like atomic congestion games, where fairness-efficiency trade-offs arise.

ACKNOWLEDGMENTS

This research was supported in part by the Simons Foundation Agency project Collaborative Research: Transferable, Hierarchical, Expressive, Optimal, Robust, Interpretable NETWORKS (THEORINET) Award MPS-MODL-00814647.

REFERENCES

- [1] Enrico Angelelli, Valentina Morandi, Martin Savelsbergh, and Maria Grazia Speranza. 2021. System optimal routing of traffic flows with user constraints using linear programming. *European journal of operational research* 293, 3 (2021), 863–879.
- [2] 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.
- [3] Soumya Basu, Ger Yang, Thanasis Lianeas, Evdokia Nikolova, and Yitao Chen. 2017. Reconciling selfish routing with social good. In *International Symposium on Algorithmic Game Theory*. Springer, 147–159.
- [4] Deeparnab Chakrabarty, Aranyak Mehta, and Viswanath Nagarajan. 2005. Fairness and optimality in congestion games. In *Proceedings of the 6th ACM Conference on Electronic Commerce*. 52–57.
- [5] George Christodoulou, Elias Koutsoupias, and Paul G Spirakis. 2011. On the performance of approximate equilibria in congestion games. *Algorithmica* 61 (2011), 116–140.
- [6] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier-Moses. 2007. Fast, Fair, and Efficient Flows in Networks. *Operations Research* 55, 2 (2007), 215–225. <https://doi.org/10.1287/opre.1070.0383> arXiv:<https://doi.org/10.1287/opre.1070.0383>
- [7] Frank A Farris. 2010. The Gini index and measures of inequality. *The American Mathematical Monthly* 117, 10 (2010), 851–864.
- [8] Bryce L. Ferguson and Jason R. 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>
- [9] Sreenivas Gollapudi, Kostas Kollias, Alkmini Sgouritsa, and Ali Kemal Sinop. 2025. Fairness and Optimality in Routing. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems*. 923–931.
- [10] Olaf Jahn, Rolf H. Möhring, Andreas S. Schulz, and Nicolás E. 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> arXiv:<https://doi.org/10.1287/opre.1040.0197>
- [11] Devansh Jalota, Kiril Solovey, Matthew Tsao, Stephen Zoepf, and Marco Pavone. 2023. Balancing fairness and efficiency in traffic routing via interpolated traffic assignment. *Autonomous Agents and Multi-Agent Systems* 37, 2 (2023), 32.
- [12] Tim Roughgarden. 2001. Stackelberg scheduling strategies. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing (Hersonissos, Greece) (STOC '01)*. Association for Computing Machinery, New York, NY, USA, 104–113. <https://doi.org/10.1145/380752.380783>
- [13] Tim Roughgarden. 2002. How unfair is optimal routing?. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, California) (SODA '02)*. Society for Industrial and Applied Mathematics, USA, 203–204.
- [14] Tim Roughgarden. 2002. The price of anarchy is independent of the network topology. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (Montreal, Quebec, Canada) (STOC '02)*. Association for Computing Machinery, New York, NY, USA, 428–437. <https://doi.org/10.1145/509907.509971>
- [15] 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>
- [16] Andreas S Schulz and Nicolás E Stier-Moses. 2006. Efficiency and fairness of system-optimal routing with user constraints. *Networks: An International Journal* 48, 4 (2006), 223–234.
- [17] Yosef Sheffi. 1985. *Urban transportation networks*. Vol. 6. Prentice-Hall, Englewood Cliffs, NJ.
- [18] Transportation Networks for Research Core Team. 2016. Transportation Networks for Research. <https://github.com/bstabler/TransportationNetworks>. Accessed October 7, 2025.