Fairness and Optimality in Routing

Sreenivas Gollapudi
Google Research
California, USA
sgollapu@google.com

Kostas Kollias
Google Research
California, USA
kostaskollias@google.com

Alkmini Sgouritsa
Athens University of Economics and Business, and
Archimedes/Athena RC
Athens, Greece
alkmini@aueb.gr

Ali Kemal Sinop
Google Research
California, USA
asinop@google.com

ABSTRACT

We study the existence of almost fair and near-optimal solutions to a routing problem as defined in the seminal work of Rosenthal [41]. We focus on the setting where multiple alternative routes are available for each potential request (which corresponds to a potential user of the network). This model captures a collection of diverse applications such as packet routing in communication networks, routing in road networks with multiple alternative routes, and the economics of transportation of goods.

Our proposed centralized routes have provable guarantees in terms of both the total cost and fairness concepts such as approximate envy-freeness. We employ and appropriately combine tools from algorithmic game theory and fair division. Our results apply on two distinct models: the splittable case where the request is split among the selected paths (e.g., routing a fleet of trucks) and the unsplittable case where the request is assigned to one of its designated paths (e.g., a single user request). Finally, we conduct an empirical analysis to test the performance of our approach against simpler baselines using the real world road network of New York City.

KEYWORDS

Congestion Games; Nash equilibrium; Envy-freeness

ACM Reference Format:
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 (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

Online route recommendation platforms are widely popular systems that help users navigate road networks by serving billions of a requests on a daily basis. In addition to providing navigation directions, these systems have the capacity to provide important dynamic information to the user, such as road segment delays, closures, etc., which in turn are used to provide high quality route recommendations. The approach that is followed by the vast majority (if not all) such systems is natural: receive a navigation request and route it optimizing for the experience of the user who sent it. However, as these systems become more and more widespread, it becomes an interesting question of how one should route requests with more general objectives in mind. One natural objective is of course to minimize the *social cost*, i.e., the aggregate travel time of all the drivers in the road network. In his seminal work, Rosenthal defined the *network congestion game* model for routing [41] which can be used for such optimization tasks. In this model each agent wishes to route a unit of (unsplittable) flow between her two endpoints and the cost of each edge e scales with the number of agents ℓ_e that use it, as given by a congestion function $c_e(\ell_e)$. Rosenthal also considered a *weighted* or *splittable* model where the demand of an agent is an integer $w > 1$ that can be split on w different paths.

Another natural objective is to guarantee fairness among the agents, for instance, it should be the case that two routing requests with the same origin and destination will not suffer a very different cost. However, one can show that this is precisely what happens in solutions that minimize the aggregate cost under various models (e.g., in the well known Pigou's example [39]). The work of Chakrabarty et al. [13] bounds this "unfairness" when the latencies (congestion functions) are linear. The problem of finding "fair" or approximately "fair" outcomes is explored in various contexts in the literature, such as nonatomic settings [18, 27, 28, 42], bandwidth allocation [29], and scheduling problems [33].

However, the interplay between cost and fairness creates a trade-off that one needs to balance, instead of focusing on one of the two objectives. Our work explicitly explores this trade-off by suggesting meaningful routing solutions with guarantees on both objectives. We further complement our results by providing a range of lower bounds on approximations of both objectives. To the best of our knowledge our work is the first to explore this trade-off in Rosenthal's atomic routing model. The only closely related work in this direction is the very recent work of Jalota et al. [28] that considers the minimization of the social cost under the constraint that the unfairness between agents that share the same origin and destination is bounded by some parameter in a *nonatomic routing* setting, in which agents are modeled as flow particles and so a convex program can be used for optimization. The model that we study is a discrete setting where each agent controls one unit of flow (or multiple units in the weighted case) which results in a computationally intractable problem. To address this we need to come up

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

with other ways to suggest (approximate) routing solutions, such as using tools provided by algorithmic game theory.

In Rosenthal's model [41], the travel time on each edge depends on the congestion on it and it is denoted by a latency/cost function. A characteristic in these settings is that often latency functions are not precisely known up front. Consider for instance the case of road networks, for which the traffic-induced delays on road segments are predicted with a margin of error. This inaccuracy can be attributed to various sources, such as a limited understanding of how much traffic will actually appear in the network or insufficient learning of delay functions from past data. However, with the widespread use of ML techniques and the availability of historical data, it is reasonable to assume access to an ML predictor that achieves a bounded error. We consider edge latency functions that are *polynomial* functions, an expressive delay function model which includes the functional form of Bureau of Public Roads [10] and has been used in various studies such as [16, 32].

In this context, we focus on Rosenthal's network routing models (unweighted and weighted) and provide theoretical and experimental insights. We first suggest routes with provable guarantees both in terms of the social cost and fairness and show how their performance degrades with errors in the edge cost predictions. Since those theoretical results are derived from a worst-case analysis, we put our approach to the test in a real road network and show that in this realistic scenario the performance of our routing solution is dramatically improved.

1.1 Our contributions

We study Rosenthal's network routing models, considering both settings: in the first one each agent has unit demand and is assigned a single path, whereas in the second one each agent i has a demand w_i and is, hence, assigned w_i (not necessarily different) paths. We focus on the case of edge latency functions that are polynomials of maximum degree d. The metrics we seek to optimize are the *social cost*, i.e., the aggregate cost of all agents in the network, and the *envy-ratio*, i.e., the maximum ratio between (average) costs of agents with the same endpoints (see Section 2 for a complete definition). To the best of our knowledge, we are the first to investigate the trade-off between those two metrics.

We first consider the unit demand/unweighted case. Our first result shows that the envy-ratio of an optimal path assignment is at most $2^{d+1} - 1$ and this is tight (**Theorem 3.1**); only the upper bound of 3 for linear edge congestion functions was known from prior work [13]. We then show that in order to achieve an envy-ratio better than 2^d, we need to incur a high social cost and violate other desirable properties (**Theorems 3.2 and 3.3**). For envy-ratio at least 2^d we first provide lower bounds on the approximation of the minimum social cost (**Theorem 3.4**). We then show that the best Nash equilibrium of the routing game guarantees an envy ratio of 2^d and a $d + 1$ approximation to the minimum social cost (**Theorem 3.5**); the approximation to the minimum social cost is due to [14]. Note that the Nash equilibrium in the nonatomic setting is a flow with envy-ratio equal to 1.

Then, we turn our attention to the more general weighted model with (integer) splittable flows. We design an algorithm that applies any positive result that we obtain in the unweighted model (e.g. in Theorems 3.1 and 3.5) to the weighted model. Specifically, given a weighted setting, the algorithm computes a desirable outcome for the unweighted model where each agent i is replaced by w_i agents and then assigns paths to the agents in a round-robin fashion from best to worst. By considering the same weight/demand w for all agents, **Theorem 4.1** shows that the worst-case approximation in terms of the social cost is preserved and the worst-case envy-ratio is improved by (approximately) a factor w, which we show that is tight for the optimal assignment (**Theorem 4.2**). Those results are then extended in the case of different weights (**Theorem 4.3**).

Finally, we conduct an experimental study to show the efficacy of our proposed solutions in a real world setting. We use higher order polynomial delay functions (as in [10, 16, 32]) and generate predictions on the travel times with a bounded error. We generate synthetic demands and calibrate them to the point where the observed route costs of selected demands match the ones observed in online routing platforms. We show that our algorithm outperforms natural baselines in terms of fairness while matching them in terms of the total cost. Then, we examine the case that the edge congestion functions are not accurate, but are rather provided by an ML predictor. In **Theorem 5.1** we bound the distortion that is caused by the prediction error to any solution with a given envy-ratio and social cost (with respect to the predicted functions).

We summarize the novelty of our work as follows:

(1) We extend the upper bound of [13] to polynomial costs and provide a tight lower bound.
(2) To the best of our knowledge, we are the first to provide theoretical upper and lower bounds with respect to the optimality-fairness tradeoff in the Rosenthal's model.
(3) We observe undesirable and unnatural outcomes are inevitable when seeking low envy-ratio.
(4) We extent our findings to the weighted case where each agent controls a fleet of vehicles.

1.2 Related Work

Our work is related to the literature on fairness in congestion games, unsplittable and splittable congestion games, and algorithms with ML predictions.

Congestion Games. The problem we consider was first introduced as a congestion game by Rosenthal [41], who studied a case with a restricted number of paths and integer flows. He showed that a pure Nash equilibrium always exists in those games by providing a potential function. Regarding the unweighted case, Meyers and Schulz [36] studied the computation of the social cost minimizing outcome. There is a line of works studying the computation [1, 5, 20] and the efficiency [2, 4, 11, 14, 15] of equilibria in those games. The nonatomic model, where each agent controls an infinitesimal amount of traffic, has also been studied [44, 45].

Haurie and Marcotte [25] initiated the study of a more general model with arbitrary (splittable) flows. Computation and existence of equilibria in those games were studied in [3, 8, 9, 23, 24, 34, 38]. The inefficiency of equilibria in this model has been also studied [17, 22, 43]. Baier et al. [7] proposed the $k-$splittable variation where each agent is allowed to use a limited number of paths to route their traffic. There is a long line of works considering the complexity and approximability of this problem (e.g. [12, 21, 31, 35, 46]).

Fairness in Congestion Games. Jalota et al. [28] studied the minimization of social cost in nonatomic setting under the constraint of bounded unfairness between agents with the same origin and destination. In the same setting, Jahn et al. [27] considered the unfairness of the paths' "normal" length (they consider path lengths that do not depend on the congestion). Chakrabarty et al. [13] studied fair allocations on the basis of minmax fairness[1] in symmetric unweighted network congestion games (among other settings). They proved that finding the minmax cost and therefore the fair allocation, when the edge latencies are linear, is NP-hard; in this case they showed that the optimal cost is a constant approximation to the fair allocation. Correa et al. [18] studied similar notions of fairness and showed similar hardness results for the nonatomic setting. Kleinberg et al. [29] considered the bandwidth allocation of unsplittable flow in single source directed graphs and gave algorithms that approximate the maxmin fair allocation. Roughgarden [42] quantified the *unfairness* in the symmetric network nonatomic setting, where the unfairness is defined as the maximum ratio between the latency of a path in an optimum outcome and the latency of a path in a Nash equilibrium. Lipton et al. [33] defined the notion of envy-ratio in scheduling problems with identical processors, as the ratio between the maximum and the minimum completion time. It is worth mentioning that fairness among agents has also been studied in other, not that related but very popular contexts, such as fair-division of goods, e.g., the famous cake-cutting problem (see [40]) and algorithmic fairness that targets on socially unbiased outcomes [19, 26].

2 PRELIMINARIES

Network congestion setting. A *network congestion problem* is described by the following tuple

$$(N, (w_i)_{i\in N}, G, (c_e)_{e\in E(G)}, (S_i)_{i\in N}),$$

where $N = \{1, \dots, n\}$ is a set of agents, w_i is a positive integer representing the weight of agent $i \in N$, G is an undirected graph with $E(G)$ being the set of its edges, c_e is a latency/cost function of the congestion on edge $e \in E(G)$ and S_i is the set of all possible choices/strategies for agent $i \in N$. We consider two models, the unweighted (unsplittable flow), where $w_i = 1$ for all $i \in N$, and the more general weighted (splittable flow) model. In both models, each agent i wants to route her traffic from a root r_i to a destination t_i in G. In the former model, each agent has a unit of traffic that needs to be routed via a single path. In the latter model, each agent i has w_i units of traffic that are routed via w_i, not necessarily different or disjoint, paths (still each unit of traffic is routed via a single path).

Strategy profile. Each agent i has a set of basic strategies which contains all the alternative paths that connect r_i with t_i and we denote them by $\mathcal{P}_i$. In the *unweighted* setting, each agent i selects a single path to route their unit of traffic and, therefore, S_i is just $\mathcal{P}_i$. In the *weighted* setting, each agent i may split her traffic and route it via many paths and therefore, $S_i = \mathcal{P}_i^{w_i}$; in other words, agent i chooses w_i *not necessarily different* paths, i.e., $s_i = \{p_{i1}, \dots, p_{iw_i}\} \in S_i$. A *strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$ is a vector of strategies where $s_i \in S_i$ is a strategy for agent i.

[1] A strategy profile is *minmax fair* if there is no way to decrease the cost of any agent i without increasing the cost of another agent j who was already experiencing a higher cost than i.

Symmetric congestion network. A congestion network is *symmetric* if all agents have the same basic strategies, i.e., for every two agents i and j, $\mathcal{P}_i = \mathcal{P}_j$. In the unweighted setting, or in the weighted setting where all agents have the same weight, all agents have the same strategy space, i.e., $S_i = S_j$, for any i, j.

Agents' latency/cost. Given a strategy profile $\mathbf{s}$, let $\ell_e(\mathbf{s})$ be the load on edge e due to the congestion caused by $\mathbf{s}$; in the unweighted setting, $\ell_e(\mathbf{s}) = |\{s_i \mid e \in s_i, i \in N\}|$ and in the weighted setting, $\ell_e(\mathbf{s}) = |\{p_{ij} \mid e \in p_{ij}, i \in N, 1 \le j \le w_i\}|$. Then, the latency on e is given by $c_e(\ell_e(\mathbf{s}))$ or simply $c_e(\mathbf{s})$ and the latency of any path p is given by $c_p(\mathbf{s}) = \sum_{e\in p} c_e(\mathbf{s})$. In the unweighted case, s_i is just a path and so, each agent i experiences a latency of $C_i(\mathbf{s}) = c_{s_i}(\mathbf{s})$. In the weighted case, s_i is a set of paths and so, each agent i experiences a latency of $C_i(\mathbf{s}) = \sum_{p\in s_i} c_p(\mathbf{s})$.

Polynomial cost functions. We consider *polynomial* latency functions that are very expressive and can capture many different scenarios regarding the dependency of the latency on the congestion. A latency function $c_e(\ell)$ is polynomial if it is of the form $c_e(\ell) = \sum_{r=0}^{d} a_{e,r}\ell^r$, for some constants $a_{e,r} \ge 0$. A special polynomial cost functions are the *affine* cost functions where $d = 1$.

We next discuss our two objectives in the network congestion problem, which are to minimize the *social cost* and to minimize the *envy-ratio*, and one constraint, which is *local Pareto-efficiency*.

Social cost. The *social cost* of a strategy profile $\mathbf{s}$ is given by $SC(\mathbf{s}) = \sum_{i\in N} C_i(\mathbf{s})$. The optimum, i.e., the social cost minimizing solution, is defined as $\min_{\mathbf{s}} SC(\mathbf{s})$ over all strategy profiles $\mathbf{s}$.

Envy-ratio. Inspired by fair division of a resource or indivisible goods we adapt fairness notions in network congestion settings and more specifically, we propose the use of the envy-ratio that have been considered in scheduling problems [33]. We remark that this is the same notion with the minimax fairness studied in [13]. In any network congestion instance let A be an outcome and $C_i(A)$ be the latency/cost for agent i under A. The envy-ratio of A is the minimum $\alpha \ge 1$ such that for any two agents i and j with $S_i = S_j$, $C_i(A) \le \alpha C_j(A)$ (in the case of agents with different weights, the cost $C_i(A)$, is replaced with the cost per unit weight $C_i(A)/w_i$).[2] We remark that the envy-ratio is equivalent to approximate envy-freeness, where A is α envy-free (α−EF) if any agent envies any other agent by at most α, i.e. for any agents i and j, $C_i(A) \le \alpha C_j(A)$.

Local Pareto-efficiency. An outcome is *Pareto-efficient* if there is no other solution where no agent is worse off (experiences more latency) and at least one agent is better off (experiences strictly less latency). Requiring a Pareto-efficient outcome is desirable but very restricted and usually difficult to find. We relax this concept and define the *local Pareto-efficiency*. An outcome is *locally Pareto-efficient* if there is no other solution derived by a *unilateral deviation*, i.e., by a single agent changing their strategy, where no agent is worse off and at least one agent is better off. We require that our outcome satisfies local Pareto-efficiency. To justify this restriction, note that an outcome doesn't satisfy local Pareto-efficiency when there are resources/edges that are not used, and by using them some agents' latency improves. Obviously, this is not a desirable

[2] Note that we only compare the latencies of agents with the same set of alternative paths (agents with the same root and destination). The reason is that otherwise the latencies may differ due to different distances and not due to unfairness. When it comes to the general case it is not clear what is the right way to define the envy-ratio and one may need to normalize according to some ground latency.

outcome and it is only natural to use any unused resource in order to improve some agents' latency without harming anybody else.

Finally, a useful concept for our results is *Nash equilibrium*.

Nash equilibrium. A (pure) *Nash equilibrium* is a strategy profile $\mathbf{s}$ such that for any agent i and any alternative strategy $s'_i \in S_i$ for agent i, $C_i(\mathbf{s}) \le C_i(s'_i, \mathbf{s}_{-i})$.

3 RECOMMENDED ROUTING AND ANALYSIS - UNWEIGHTED MODEL

In this section we investigate the trade-off between minimizing the social cost and minimizing the envy in the case where every agent has a unit weight (unweighted model). We provide the exact envy-ratio for the minimum social cost outcome, and inapproximability bounds as we move away from the minimum social cost. We complement our results by showing the existence of a locally Pareto-efficient outcome with the minimum envy-ratio possible; this outcome is any pure Nash equilibrium. The price of stability bound known in the literature [14] provides the best approximation of such an outcome to the minimum social cost. Finally, we consider the case where the cost functions may be inaccurate and their coefficients are given by an ML prediction.

3.1 Envy-Ratio in the Minimum Social Cost Outcome

We begin our analysis by understanding what is the worst-case envy-ratio in an optimal solution in terms of the social cost. We note that a solution minimizing the social cost by definition satisfies local Pareto-efficiency. Chakrabarty et al. [13] showed that the envy-ratio of the optimal allocation is at most 3 for linear cost functions, i.e. $c_e(x) = a_e x$. Our result is a generalization to polynomial cost functions and we additionally provide a tight lower bound.

Theorem 3.1. *The envy-ratio in the optimal allocation is at most $2^{d+1} - 1$ ($d \ge 1$) and this is asymptotically tight.*

Proof. Suppose on the contrary that in the optimum $\mathbf{s}$ there are two agents i, j, with the same root and destination, such that $C_i(\mathbf{s}) > (2^{d+1} - 1)C_j(\mathbf{s})$. We show next that if agent i deviates to s_j results in a smaller social cost, contradicting $\mathbf{s}$ being the optimum.

Let $\mathbf{s}' = (s_j, \mathbf{s}_{-i})$ be the outcome derived by $\mathbf{s}$ after agent i deviates to s_j. Note that the latency doesn't change on the edges that belong either to both s_i and s_j or to neither of them. Therefore,

$$SC(\mathbf{s}') - SC(\mathbf{s}) = \sum_{e \in (s_j \setminus s_i) \cup (s_i \setminus s_j)} \left(\ell_e(\mathbf{s}')c_e(\mathbf{s}') - \ell_e(\mathbf{s})c_e(\mathbf{s})\right). \quad (1)$$

Consider any edge $e \in s_j \setminus s_i$ and let $\ell = \ell_e(\mathbf{s})$; then $\ell_e(\mathbf{s}') = \ell + 1$. Next, for simplicity we drop the index e from the coefficients. The difference of the right-hand side of (1) for e becomes,

$$\begin{aligned}
& \ell_e(\mathbf{s}')c_e(\mathbf{s}') - \ell_e(\mathbf{s})c_e(\mathbf{s}) \\
= \; & (\ell+1)\sum_{r=0}^{d} a_r(\ell+1)^r - \ell \sum_{r=0}^{d} a_r \ell^r \\
= \; & \ell \sum_{r=0}^{d} a_r \sum_{k=0}^{r} \binom{r}{k} \ell^k - \ell \sum_{r=0}^{d} a_r \ell^r + \sum_{r=0}^{d} a_r(\ell+1)^r \\
= \; & \sum_{r=1}^{d} a_r \sum_{k=0}^{r-1} \binom{r}{k} \ell^{k+1} + \sum_{r=0}^{d} a_r(\ell+1)^r \\
\le \; & \sum_{r=1}^{d} a_r \left(\ell^r \sum_{k=0}^{r-1} \binom{r}{k} + (2\ell)^r\right) \\
= \; & \sum_{r=1}^{d} a_r \ell^r \left(2^r - 1 + 2^r\right) \le (2^{d+1} - 1)c_e(\mathbf{s}),
\end{aligned}$$

where the first inequality comes from the fact that $\ell \ge 1$, because agent j is using e. Consider now any edge $e \in s_i \setminus s_j$ and again let $\ell = \ell_e(\mathbf{s})$; then $\ell_e(\mathbf{s}') = \ell - 1$. We again drop the index e from the coefficients of the polynomial cost function. The difference of the right-hand side of (1) for e becomes,

$$\begin{aligned}
& \ell_e(\mathbf{s}')c_e(\mathbf{s}') - \ell_e(\mathbf{s})c_e(\mathbf{s}) \\
= \; & (\ell-1)\sum_{r=0}^{d} a_r(\ell-1)^r - \ell \sum_{r=0}^{d} a_r \ell^r \\
= \; & \ell \sum_{r=0}^{d} a_r \sum_{k=0}^{r} \binom{r}{k} \ell^k (-1)^{r-k} - \ell \sum_{r=0}^{d} a_r \ell^r - \sum_{r=0}^{d} a_r(\ell-1)^r \\
= \; & \sum_{r=1}^{d} a_r \sum_{k=0}^{r-1} \binom{r}{k} \ell^{k+1}(-1)^{r-k} - \sum_{r=0}^{d} a_r(\ell-1)^r \\
\le \; & \sum_{r=1}^{d} a_r \ell^r \left(\sum_{k=0}^{r} \binom{r}{r-k}(-1)^{r-k} - 1\right) - a_0 \\
= \; & -\sum_{r=1}^{d} a_r \ell^r - a_0 = -c_e(\mathbf{s}),
\end{aligned}$$

where for the second to last equality we used the fact that for $r > 0$, $\sum_{k=0}^{r} \binom{r}{r-k}(-1)^{r-k} = \sum_{k'=0}^{r} \binom{r}{k'}(-1)^{k'} = 0$. Using the above inequalities in (1), we get

$$\begin{aligned}
SC(\mathbf{s}') - SC(\mathbf{s}) \; & \le \; (2^{d+1} - 1) \sum_{e \in s_j \setminus s_i} c_e(\mathbf{s}) - \sum_{e \in s_i \setminus s_j} c_e(\mathbf{s}) \\
& \le \; (2^{d+1} - 1) \sum_{e \in s_j} c_e(\mathbf{s}) - \sum_{e \in s_i} c_e(\mathbf{s}) \\
& = \; (2^{d+1} - 1)C_j(\mathbf{s}) - C_i(\mathbf{s}) < 0,
\end{aligned}$$

where the second inequality comes after adding the non-negative term $(2^{d+1} - 1)\sum_{e \in s_j \cap s_i} c_e(\mathbf{s}) - \sum_{e \in s_j \cap s_i} c_e(\mathbf{s})$ for $d \ge 0$. The last inequality comes from our assumption that the envy-ratio is more than $2^{d+1} - 1$. $SC(\mathbf{s}') < SC(\mathbf{s})$ means that $\mathbf{s}$ is not optimal which is a contradiction. Therefore, the envy-ratio is at most $2^{d+1} - 1$.

This upper bound is asymptotically tight as we show next. Consider the network in Figure 1 with two agents where both have root r and destination t. There are only two edges connecting r to t with costs functions $x^d + 2^{d+1} - 2 - \varepsilon$, for $\varepsilon > 0$, and x^d, respectively. It is easy to verify that the optimum solution is to route the agents via different paths. This results in envy-ratio of $2^{d+1} - 1 - \varepsilon$ that converges to the upper bound as ε goes to zero. $\square$

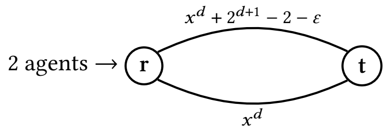

Figure 1: Lower bound on the envy-ratio of the optimum solution in the unweighted model.

3.2 Envy-Ratio and Social Cost Trade-off

In this subsection we investigate the trade-off between the envy-ratio and social cost objectives. We first present a sequence of impossibility results, which indicate that for a small envy-ratio we must have a high social cost and also violate local Pareto-efficiency.

Theorem 3.2. *For any $1 < \beta \leq 2^d$, we cannot guarantee both envy-ratio less than β and approximation ratio to the optimum social cost less than $(q+1)^d - \varepsilon$, where $q = \lfloor \frac{1}{\beta^{1/d}-1} \rfloor$ and $\varepsilon > 0$ is an arbitrarily small positive value.*

Proof. Consider the network of Figure 2, where $n = k \cdot q! + 1$, for some positive integer k, agents have root r and destination t and there are $n-1$ edges of the same cost function x^d.

First notice that in any outcome where there exists an edge with $\ell \leq q$ agents, it cannot be that all used edges have ℓ agents and therefore the envy-ratio would be at least $\frac{(q+1)^d}{q^d} = (1 + \frac{1}{q})^d \geq (1 + \beta^{1/d} - 1)^d = \beta$. So, if the envy-ratio is less than β, all agents have cost at least $(q+1)^d$.

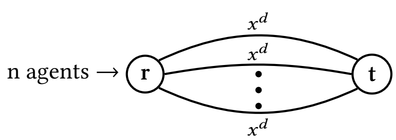

Figure 2: Trade-off between the envy-ratio and the approximation ratio, for $\beta \leq 2^d$.

The minimum social cost appears when one edge is used by two agents and all other edges are used by exactly one agent and it is $n - 2 + 2^{d+1}$. So, the approximation ratio is at least $\frac{n(q+1)^d}{n-2+2^{d+1}} = (q+1)^d - \frac{(2^{d+1}-2)(q+1)^d}{n+2^{d+1}-2}$; note that n can be arbitrarily high by increasing the parameter k and therefore, there exists sufficiently large n such that $\frac{(2^{d+1}-2)(q+1)^d}{n+2^{d+1}-2} \leq \varepsilon$. Then, the theorem follows. □

Theorem 3.3. *For any $1 < \beta \leq 2^d$, we cannot guarantee both envy-ratio less than β and satisfying local Pareto-efficiency.*

Proof. Consider the example of Figure 2 with n agents and $n-1$ edges. In the optimum, one edge is used by two agents and all other edges are used by exactly one agent. This outcome has envy-ratio 2^d. In any other outcome, there are unused edges and edges that are used by at least 2 agents. Consider some agent i that uses one of the latter edges. i's cost is at least 2^d and with a unilateral deviation to an empty edge, it would drop to 1. Therefore, we cannot have an outcome that is locally Pareto-efficient and has envy-ratio less than 2^d in this instance. □

Theorem 3.4. *For any $2^d < \beta \leq 2^{d+1} - 1$, we cannot guarantee both envy-ratio less than β and approximation ratio to the optimum social cost less than $\frac{2^{d+1}}{1+\beta}$.*

Proof. To show this we consider the network of Figure 3 with two agents with root r and destination t. If the agents use different paths, the envy-ratio is β. The only way that the envy-ratio is less than β is when both agents use the same edge. The solution where both agents use the same edge, that has the minimum social cost, is when both agents use the lower edge with social cost $2 \cdot 2^d = 2^{d+1}$. The optimal social cost appears when the agents use different edges and it is $1 + \beta$. Therefore, the best approximation ratio we can have for envy-ratio less than β is $\frac{2^{d+1}}{1+\beta}$. □

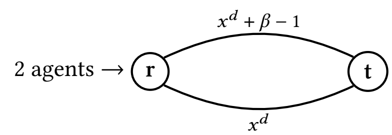

Figure 3: Trade-off between the envy-ratio and the approximation ratio for $\beta > 2^d$.

3.3 Locally Pareto-Efficient Outcome with the Minimum Envy-Ratio

We now present a positive result: there exists an outcome, namely the pure Nash equilibrium, that guarantees an envy-ratio at most 2^d. A pure Nash equilibrium always exists [41] and the inefficiency of the best equilibrium in terms of the social cost is at most $d+1$ [14]. Note that any Nash equilibrium is trivially locally Pareto-efficient, and therefore this is the best outcome in terms of envy-ratio that is locally Pareto-efficient. We then get the following theorem.

Theorem 3.5. *There exists an outcome that is locally Pareto-efficient, has envy-ratio at most 2^d, and its approximation ratio against the minimum social cost is at most $d+1$.*

Proof. Suppose that $\mathbf{s}$ is any pure Nash equilibrium in the network congestion game. This means that no player may decrease their latency by unilaterally deviating to another strategy/path.

We first show that the envy-ratio in $\mathbf{s}$ (i.e., in *any* pure Nash equilibrium) is at most 2^d. For this consider any two agents i and j with the same root and destination and let $\mathbf{s}' = (s_j, \mathbf{s}_{-i})$ be the outcome derived by $\mathbf{s}$ after agent i deviates to s_j. The fact then $\mathbf{s}$ is a pure Nash equilibrium means that agent i cannot improve his cost by choosing s_j and therefore, $C_i(\mathbf{s}) \leq C_i(\mathbf{s}')$.

Note that when i deviates from s_i to s_j, the load doesn't change on the edges that belong to both s_j and s_i and increases by 1 on the edges belonging to s_j but not to s_i. So,

$$\begin{aligned} C_i(\mathbf{s}) \quad &\leq \quad C_i(\mathbf{s}') = \sum_{e \in s_j \setminus s_i} c_e(\mathbf{s}') + \sum_{e \in s_j \cap s_i} c_e(\mathbf{s}') \\ &= \quad \sum_{e \in s_j \setminus s_i} \sum_{r=0}^{d} a_{e,r} (\ell_e(\mathbf{s}) + 1)^d + \sum_{e \in s_j \cap s_i} c_e(\mathbf{s}) \end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{e \in s_j \setminus s_i} \sum_{r=0}^{d} a_{e,r} (2\ell_e(\mathbf{s}))^d + \sum_{e \in s_j \cap s_i} c_e(\mathbf{s}) \\
&\leq 2^d \sum_{e \in s_j \setminus s_i} c_e(\mathbf{s}) + \sum_{e \in s_j \cap s_i} c_e(\mathbf{s}) \leq 2^d C_j(\mathbf{s}),
\end{aligned}
$$

where the second inequality comes from the fact that for any edge $e \in s_j$, $\ell_e(\mathbf{s}) \geq 1$, since agent j is using e.

The upper bound on the approximation ratio to the minimum social cost comes from known bounds on the price of stability which is the ratio between the social cost of the best pure Nash equilibrium[3] and the minimum social cost. So, if we consider a pure Nash equilibrium that minimizes the social cost, the price of stability provides the approximation ratio to the minimum social cost. The exact price of stability for general (not only network) congestion games with polynomial cost functions is given in [14] and it is the following:

$$
\max_{r>1} \frac{(2^d d + 2^d - 1) \cdot r^{d+1} - (d+1) \cdot r^d + 1}{(2^d + d - 1) \cdot r^{d+1} - (d+1) \cdot r^d + 2^d d - d + 1}, \qquad (2)
$$

which is upper bounded by $d+1$ and converges to this value as d increases [14]. □

Corollary 3.6. *For affine cost functions, i.e., for $d = 1$, there exists an outcome (namely the best pure Nash equilibrium) with envy-ratio at most 2 and approximation ratio at most 1.577. Additionally, there is no outcome with envy-ratio at most 2 and approximation ratio better than 4/3.*

Proof. From Theorem 3.5, if we set $d = 1$ we get that any pure Nash equilibrium admits an envy-ratio of at most 2. By setting again $d = 1$ in (2), the best pure Nash equilibrium has an approximation ratio to the minimum social cost of at most 1.577. The rest of the results come from Theorem 3.4 for $d = 1$. □

Remark 1. *The bound on the approximation ratio from [14] is tight for general congestion games but provide only an upper bound for the network congestion games. It may be the case that this bound may improve for network congestion games or that our lower bound in Theorem 3.4 is not tight. This means that it is still possible that the best pure Nash equilibrium is the outcome with envy-ratio at most 2^d and the best approximation ratio against the minimum social cost.*

4 RECOMMENDED ROUTING AND ANALYSIS - WEIGHTED MODEL

We now study the weighted model. Our main contribution is an algorithm that allocates paths to the agents, and a corresponding theorem for the theoretical guarantees on the induced solution in terms of the social cost and the envy-ratio. The theorem reduces the guarantees of the unweighted model to guarantees for the weighted model. We first give our result under the assumption that all weights are the same, i.e. $w_i = w$ for all agents i, where w is a global integer parameter. Then, we extend our result to the more complicated case of different weights.

[3]By best pure Nash equilibrium we mean the pure Nash equilibrium with the minimum social cost.

4.1 Weighted Model with Same Weights

In this subsection we assume that the agents have all the same weight. We design a simple algorithm that solves a problem with more agents with unit weights and allocates the paths of this solution to the weighted agents in a round-robin fashion.

Theorem 4.1. *Given a weighted congestion setting denoted by $\mathcal{G} = (N, (w)_{i \in N}, G, (c_e)_{e \in E(G)}, (S_i)_{i \in N})$, let N^w be the set of agents derived after replacing each agent of N with w agents. Suppose that there exists an outcome $\mathbf{s}'$ of the unweighted setting $\mathcal{G}' = (N^w, (1)_{i \in N^w}, G, (c_e)_{e \in E(G)}, (\mathcal{P}_i)_{i \in N^w})$, where the envy-ratio is at most β and the approximation ratio to the minimum social cost is at most α. Then, for $\mathcal{G}$, there exists an outcome $\mathbf{s}$ with envy-ratio at most $1 + \frac{\beta - 1}{w}$ and approximation ratio at most α.*

Proof. Suppose that in $\mathcal{G}$ there are n_{rt} agents with root r and destination t. Considering $\mathbf{s}'$, there are $n_{rt}w$ (not necessarily different) paths from r to t, each one selected by each agent of the unweighted model with root r and destination t. We will construct $\mathbf{s}$ by carefully allocating w of those paths to each of the n_{rt} agents of the weighted model. We do this in a round-robin fashion after ordering the cost of the paths in a non-decreasing order. We give Algorithm 1 that assigns w, not necessarily different, paths to each agent of the weighted case.

Algorithm 1: Round Robin on Weighted Model with equal weights w

$N^w \leftarrow$ Set of agents derived after replacing each agent of N with w agents
$\mathbf{s}' \leftarrow$ Outcome on $(N^w, (1)_{i \in N^w}, G, (c_e)_{e \in E(G)}, (\mathcal{P}_i)_{i \in N^w})$ with envy-ratio $\leq \beta$ and approx. ratio $\leq \alpha$
$\mathbf{s}' \leftarrow$ Sort $\mathbf{s}'$ by non decreasing latency
$\mathbf{s} \leftarrow |N|$ empty sets of paths
$i \leftarrow 0$
for p *in* $\mathbf{s}'$ **do**
  $\mathbf{s}[i]$.Insert(p)
  $i \leftarrow (i+1) \mod (n)$
end
Return $\mathbf{s}$

First note that $\mathbf{s}$ and $\mathbf{s}'$ have exactly the same social cost, because the same amount of traffic passes through each edge. Additionally, the two instances have same demand (i.e., same units of traffic to be routed from any vertex r to any vertex t), meaning that they have the same minimum social cost. Hence, since the approximation ratio of $\mathbf{s}'$ is α, it holds that the approximation ratio of $\mathbf{s}$ is also α.

We next show the upper bound on the envy-ratio. For each agent i of the weighted setting, let $s_i = \{p_{i1}, \ldots, p_{iw_i}\}$, where the paths are ordered according to the order that Algorithm 1 allocates them to i. The next two key claims compare the cost of the allocated paths between any two agents.

Claim 1. *For any two agents i and j in the weighted setting, $\sum_{p \in s_i \setminus \{p_{iw}\}} c_p(\mathbf{s}) \leq \sum_{p \in s_j \setminus \{p_{j1}\}} c_p(\mathbf{s})$.*

Proof. Due to the way we assigned the paths in the round-robin fashion, it is $c_{p_{ik}}(\mathbf{s}) \leq c_{p_{j(k+1)}}(\mathbf{s})$ for every $k \in \{1, \ldots, w-1\}$. After summing over all $k \in \{1, \ldots, w-1\}$, the claim follows. □

Claim 2. *For any two paths p_i and p_j that are used under* $\mathbf{s}$, $c_{p_i}(\mathbf{s}) \leq \beta \cdot c_{p_j}(\mathbf{s})$.

Proof. According to Algorithm 1, there exist agents i' and j' of the unweighted case, such that $s'_{i'} = p_i$ and $s'_{j'} = p_j$. Since $\mathbf{s}'$ has envy-ratio at most β, $c_{s_i}(\mathbf{s}') \leq \beta \cdot c_{s_j}(\mathbf{s}')$, which equivalently means that $c_{p_i}(\mathbf{s}) \leq \beta \cdot c_{p_j}(\mathbf{s})$ (recall that $\mathbf{s}$ and $\mathbf{s}'$ result in the same latency at each edge). □

By using Claims 1 and 2 we get the envy-ratio upper bound as we show next. For any agents i and j in the weighted setting,

$$\begin{aligned} C_i(\mathbf{s}) &= \sum_{p \in s_i} c_p(\mathbf{s}) \leq \sum_{p \in s_j \setminus \{p_{j1}\}} c_p(\mathbf{s}) + c_{p_{iw}}(\mathbf{s}) \\ &\leq \sum_{p \in s_j \setminus \{p_{j1}\}} c_p(\mathbf{s}) + \beta \cdot c_{p_{j1}}(\mathbf{s}) = \sum_{p \in s_j} c_p(\mathbf{s}) + (\beta - 1) c_{p_{j1}}(\mathbf{s}) \\ &\leq \left(1 + \frac{\beta - 1}{w}\right) \sum_{p \in s_j} c_p(\mathbf{s}) = \left(1 + \frac{\beta - 1}{w}\right) C_j(\mathbf{s}), \end{aligned}$$

where the last inequality comes from the fact that $c_{p_{j1}}(\mathbf{s}) \leq c_{p_{jk}}(\mathbf{s})$ for all $k \in \{1, \ldots, w\}$. This completes the proof of Theorem 4.1. □

The next theorem (whose proof is in the full version of the paper), Theorem 4.2, provides asymptotically tight upper bound on the envy-ratio for the optimal outcome.

Theorem 4.2. *In the weighted model with equal weights w, the best envy-ratio of any optimal outcome is exactly* $1 + \frac{2^{d+1}-2}{w}$ $(d \geq 1)$.

4.2 Weighted Model with General Weights

Here, we consider the general case where the weights can be different. For the envy-ratio it is not meaningful to compare the actual costs of the agents, because they route different amount of traffic. Instead, we use the average cost per unit for each agent. More formally, for an outcome $\mathbf{s}$, the average cost per unit for agent i with weight w_i is $\frac{C_i(\mathbf{s})}{w_i}$.

We next give a theorem similar to Theorem 4.1 (whose proof is in the full version of the paper), and for this we define ζ as a function of β and the vector of weights $\mathbf{w} = (w_1, \ldots, w_n)$ which is given by $\zeta(\beta, \mathbf{w}) = \max_{w_i, w_j} \left(\left\lceil \frac{w_i+1}{w_j+1} \right\rceil \frac{w_j+\beta}{w_i} \right)$.

Theorem 4.3. *Given a weighted congestion setting*

$$\mathcal{G} = (N, \mathbf{w}, G, (c_e)_{e \in E(G)}, (S_i)_{i \in N}),$$

let N' be the set of agents derived after replacing each agent $i \in N$ with w_i agents. Suppose that there exists an outcome $\mathbf{s}'$ of the unweighted setting $\mathcal{G}' = (N', (1)_{i \in N'}, G, (c_e)_{e \in E(G)}, (\mathcal{P}_i)_{i \in N'})$, where the envy-ratio is at most β and the approximation ratio to the minimum social cost is at most α. Then, for $\mathcal{G}$, there exists an outcome $\mathbf{s}$ with envy-ratio at most $\zeta(\beta, \mathbf{w})$ and approximation ratio at most α.

Note that the result of Theorem 4.3, when all weights are equal, does not fully match the result of Theorem 4.1. We remark that, if the ratio w_j+1/w_i+1 is an integer[4] for all the cases where $w_j \geq w_i$, then a slightly tighter analysis is possible resulting to the bound of Theorem 4.1 for equal weights.

5 EXPERIMENTAL EVALUATION

In this section we put our algorithm to the test against simpler baselines, using the road network of New York City. We remark that our theoretical results seem quite pessimistic as they consider a worst case analysis. The purpose of our experiments is to exhibit that for real networks we observe a small change in the social cost and a simultaneous more significant improvement in the envy ratio. Hence, the experiments showcase the practical implications of our algorithm.

We consider equal weights w, and for the case of $w = 1$, we compute a Nash equilibrium using best-response dynamics. For the case of $w > 1$, we additionally apply Algorithm 1. We extract the road network from OpenStreetMap [37]. We assign cost functions to the edges using the functional form of the Bureau of Public Roads [10], exactly as in [16, 32]. Specifically we set:

$$c_e(x) = \frac{\gamma t_e^f}{\beta_e^4} x^4 + t_e^f,$$

where t_e^f is the time needed to cross the edge when the road is empty, i.e., the free-flow travel time, and β_e is the capacity of the street, defined as the number of lanes multiplied by the free-flow speed. We tune parameter γ so that the induced travel times at equilibrium approximately match the ones observed in online navigation systems.

The baselines we compare against greedily assign routes to the agents and are called *Greedy* and *Marginal-Greedy* (for the analysis of such algorithms see [30] and the papers cited therein). Both routes rely on greedily assigning shortest paths to agents, however each one uses a different cost structure for the edges. In more detail, the structure for both algorithms is as follows: Begin with an empty network and assign w paths to each agent iteratively. Assign w times the shortest path to the agent under consideration, each time updating the edge costs based on the previous assignments. The cost used by *Greedy* is simply the cost that an additional atom would experience on the edge given its current congestion, i.e., with ℓ the load on edge e, the edge cost used by the algorithm is $c_e(\ell + 1/w)$. The cost used by *Marginal Greedy* is the marginal increase that an additional atom would cause to the *social cost of the edge*, i.e., by using the same notation, the cost is: $(\ell+1/w) \cdot c_e(\ell+1/w) - \ell \cdot c_e(\ell)$.

The Greedy baseline is natural in the sense that it simulates the algorithm used in actual online navigation systems: receive a request and greedily route it on the shortest available path. The Marginal-Greedy baseline is a direct extension that takes the social cost (one of our objectives) into account when routing requests one by one.

We divide the maps of New York in 8 regions and generate a demand arrival rate of 50 between each pair. To simplify computation, we generate a set of 10 candidate routes for each origin-destination pair using the well known *penalty method* [6]. It is expected that these 10 routes will cover most, if not all, ways that drivers typically

[4]This happens when all weights are equal. Another example is when each w_i is of the form $2^{k_i} - 1$, for some positive integer k_i.

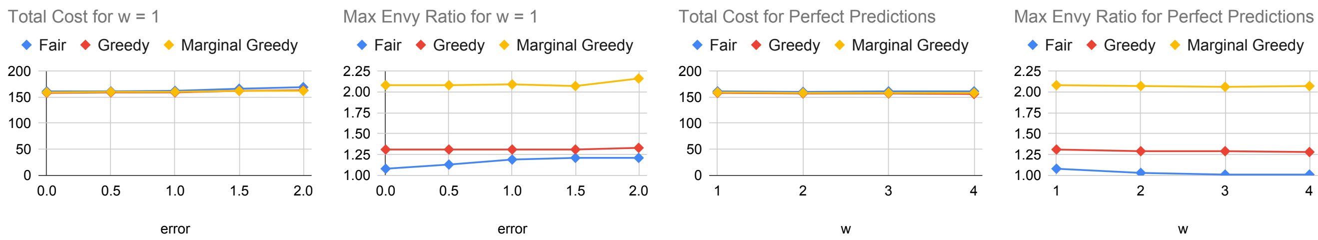

Figure 4: Our algorithm matches the total cost and improves fairness over the baselines.

traverse between the two endpoints. The strategies of players are then defined on these 10 candidate routes. We provide our instances (in the form of candidate routes and cost functions) and the code for our algorithm and baselines in our supplemental material.

We repeat our computations for $w = 1, 2, 3, 4$ on delay functions given by potentially erroneous predictions (we remind the reader that the case of $w = 1$ corresponds to some Nash equilibrium). Specifically, we generate a prediction for each $c_e(x_e)$ that is at most $(1 + \delta) \cdot c_e(x_e)$, where δ is the max error parameter. We make the following observations illustrated in Figure 4:

- There are small perturbations in the total cost across algorithms and runs. Overall the greedy baseline is performing the best on this metric but with a very small margin over the other two (typically within $1\% - 2\%$).
- Our algorithm performs significantly better than both baselines in terms of the envy ratio. For perfect predictions, an extra 3% cost appears for the most envious agent for $w = 1$ that quickly drops to near 0% as w increases. For the *Greedy* algorithm this extra cost is 33% and for the *Marginal Greedy* it is 116%.
- The envy-ratio of our algorithm decreases as we increase the number of paths w, which matches the intuition from our theoretical bounds.

5.1 Inaccurate Cost Functions

We close the section by presenting theoretical results for the case when the delay estimates are given with bounded error, specifically, the case when cost functions may be inaccurate and their coefficients are given by an ML prediction. The next theorem quantifies how this inaccuracy may impact the envy-ratio and the social cost approximation ratio of any outcome. The prediction error we consider is described as follows: For any outcome $\mathbf{s}$, let $\hat{c}_e(\mathbf{s})$ and $c_e(\mathbf{s})$ be the predicted and the actual, respectively, latency on edge e. The prediction error is the minimum values $\delta_1, \delta_2 \geq 1$ such that for each edge e and any $\mathbf{s}$, $c_e(\mathbf{s})/\delta_1 \leq \hat{c}_e(\mathbf{s}) \leq \delta_2 c_e(\mathbf{s})$.

Theorem 5.1. *Let* $\mathbf{s}$ *be any outcome with envy-ratio* β *and approximation ratio to the minimum social cost* α *according to the predicted cost functions. Then, according to the actual latency on the edges, the envy-ratio is at most* $\delta_1\delta_2\beta$ *and the approximation ratio to the minimum social cost is at most* $\delta_1\delta_2\alpha$.

The proof is given in the full version of the paper.

Remark 2. *Algorithms with predictions are often evaluated in terms of* robustness, *which implies that the best known worst case guarantees, when no predictions are provided, are recovered when the predictions are arbitrarily bad. In our settings, the worst case guarantees when the costs are unknown (which is the situation with no predictions) may be arbitrarily bad; in this regard, any algorithm, including ours, satisfies robustness. To see that consider two parallel paths between an origin and destination (e.g., a case of two islands connected by two bridges). Any algorithm will have to use one of the bridges but that bridge could have a very high cost compared to the other one (e.g., be closed altogether) incurring an arbitrarily bad social cost. The same holds with respect to the envy-ratio unless the routing algorithm always uses a single path which is a trivial case.*

6 DISCUSSION AND FUTURE DIRECTIONS

We study the traditional routing models of Rosenthal which have received significant interest historically and have many interesting application in the modern routing space, including online maps navigation, Internet packet routing, and vehicle fleet routing (trucks, ride-sharing, etc). In this work we consider one further important aspect for solving this problem beyond simply minimizing the cost: addressing fairness among agents in the network. We study the trade-off between the two objectives and give tight bounds on the envy-ratio when optimizing the social cost. We further prove that there are known notions (Nash equilibra) that optimize the envy-ratio in a meaningful way (by satisfying local Pareto-efficiency) and design algorithms with strong theoretical guarantees in terms of the twofold objective. Finally, we show that in a real road network our algorithm provides strong gains over natural baselines.

We leave as a future direction the cases where the envy-ratio is at most β for $2^d < \beta < 2^{d+1}$. Our conjecture is that there are outcomes in between the ones with the optimal social cost and the best reasonable envy-ratio. We propose the study of the best (in terms of social cost) $\alpha-$approximate Nash equilibria; those solutions loosen the condition of the pure Nash equilibrium by allowing situation where agents may be benefited by deviating but not by much. Clearly, $\alpha-$approximate Nash equilibria widen the class of outcomes comparing to Nash equilibria and it is possible that the best of them results in a better approximation ratio.

Finally, one important aspect to tackle is the computational complexity of the suggested outcomes that we discuss in the full version of the paper.

ACKNOWLEDGMENTS

This work was done when the author Alkmini Sgouritsa was a visiting researcher at Google Research, CA, USA.

REFERENCES

[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. 2008. On the impact of combinatorial structure on congestion games. *J. ACM* 55, 6 (2008), 25:1–25:22.

[2] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. 2011. Exact Price of Anarchy for Polynomial Congestion Games. *SIAM J. Comput.* 40, 5 (2011), 1211–1233.

[3] Eitan Altman, Tamer Basar, Tania Jiménez, and Nahum Shimkin. 2002. Competitive routing in networks with polynomial costs. *IEEE Trans. Autom. Control.* 47, 1 (2002), 92–96.

[4] Baruch Awerbuch, Yossi Azar, and Amir Epstein. 2005. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, Harold N. Gabow and Ronald Fagin (Eds.). ACM, 57–66.

[5] Yakov Babichenko and Aviad Rubinstein. 2021. Settling the complexity of Nash equilibrium in congestion games. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.). ACM, 1426–1437.

[6] Roland Bader, Jonathan Dees, Robert Geisberger, and Peter Sanders. 2011. Alternative Route Graphs in Road Networks. In *Theory and Practice of Algorithms in (Computer) Systems - First International ICST Conference, TAPAS 2011, Rome, Italy, April 18-20, 2011. Proceedings.* 21–32.

[7] Georg Baier, Ekkehard Köhler, and Martin Skutella. 2005. The k-Splittable Flow Problem. *Algorithmica* 42, 3-4 (2005), 231–248.

[8] Umang Bhaskar, Lisa Fleischer, Darrell Hoy, and Chien-Chung Huang. 2015. On the Uniqueness of Equilibrium in Atomic Splittable Routing Games. *Math. Oper. Res.* 40, 3 (2015), 634–654.

[9] Umang Bhaskar and Phani Raj Lolakapuri. 2018. Equilibrium Computation in Atomic Splittable Routing Games. In *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland (LIPIcs, Vol. 112)*, Yossi Azar, Hannah Bast, and Grzegorz Herman (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 58:1–58:14.

[10] Bureau of Public Roads. 1964. *Traffic assignment manual.* US Department of Commerce.

[11] Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. 2011. Tight Bounds for Selfish and Greedy Load Balancing. *Algorithmica* 61, 3 (2011), 606–637. https://doi.org/10.1007/s00453-010-9427-8

[12] Massimiliano Caramia and Antonino Sgalambro. 2008. On the approximation of the single source k-splittable flow problem. *J. Discrete Algorithms* 6, 2 (2008), 277–289.

[13] Deeparnab Chakrabarty, Aranyak Mehta, Viswanath Nagarajan, and Vijay Vazirani. 2005. Fairness and optimality in congestion games. In *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005), Vancouver, BC, Canada, June 5-8, 2005*, John Riedl, Michael J. Kearns, and Michael K. Reiter (Eds.). ACM, 52–57.

[14] George Christodoulou and Martin Gairing. 2016. Price of Stability in Polynomial Congestion Games. *ACM Trans. Economics and Comput.* 4, 2 (2016), 10:1–10:17. https://doi.org/10.1145/2841229

[15] George Christodoulou and Elias Koutsoupias. 2005. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, Harold N. Gabow and Ronald Fagin (Eds.). ACM, 67–73.

[16] Serdar Colak, Antonio Lima, and Marta Gonzalez. 2016. Understanding congested travel in urban areas. *Nature Communications* 7, 10793 (2016).

[17] Roberto Cominetti, José R. Correa, and Nicolás E. Stier Moses. 2009. The Impact of Oligopolistic Competition in Networks. *Oper. Res.* 57, 6 (2009), 1421–1437.

[18] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier Moses. 2007. Fast, Fair, and Efficient Flows in Networks. *Oper. Res.* 55, 2 (2007), 215–225.

[19] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, Shafi Goldwasser (Ed.). ACM, 214–226.

[20] Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. 2004. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, László Babai (Ed.). ACM, 604–612.

[21] Mette Gamst and B. Petersen. 2012. Comparing branch-and-price algorithms for the Multi-Commodity k-splittable Maximum Flow Problem. *Eur. J. Oper. Res.* 217, 2 (2012), 278–286.

[22] Tobias Harks. 2011. Stackelberg Strategies and Collusion in Network Games with Splittable Flow. *Theory Comput. Syst.* 48, 4 (2011), 781–802.

[23] Tobias Harks and Veerle Timmermans. 2017. Equilibrium Computation in Atomic Splittable Singleton Congestion Games. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10328)*, Friedrich Eisenbrand and Jochen Könemann (Eds.). Springer, 442–454.

[24] Tobias Harks and Veerle Timmermans. 2018. Uniqueness of equilibria in atomic splittable polymatroid congestion games. *J. Comb. Optim.* 36, 3 (2018), 812–830.

[25] Alain Haurie and Patrice Marcotte. 1985. On the relationship between Nash - Cournot and Wardrop equilibria. *Networks* 15, 3 (1985), 295–308.

[26] Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. 2019. The Disparate Effects of Strategic Manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* 2019, Atlanta, GA, USA, January 29-31, 2019*, danah boyd and Jamie H. Morgenstern (Eds.). ACM, 259–268.

[27] 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. *Oper. Res.* 53, 4 (2005), 600–616.

[28] Devansh Jalota, Kiril Solovey, Matthew Tsao, Stephen Zoepf, and Marco Pavone. 2022. Balancing Fairness and Efficiency in Traffic Routing via Interpolated Traffic Assignment. In *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor (Eds.). International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 678–686.

[29] Jon M. Kleinberg, Yuval Rabani, and Éva Tardos. 2001. Fairness in Routing and Load Balancing. *J. Comput. Syst. Sci.* 63, 1 (2001), 2–20.

[30] Max Klimm, Daniel Schmand, and Andreas Tönnis. 2019. The Online Best Reply Algorithm for Resource Allocation Problems. In *Algorithmic Game Theory - 12th International Symposium, SAGT 2019, Athens, Greece, September 30 - October 3, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11801)*, Dimitris Fotakis and Evangelos Markakis (Eds.). Springer, 200–215. https://doi.org/10.1007/978-3-030-30473-7_14

[31] Ronald Koch and Ines Spenke. 2006. Complexity and approximability of k-splittable flows. *Theor. Comput. Sci.* 369, 1-3 (2006), 338–347.

[32] Kostas Kollias, Arun Chandrashekharapuram, Lisa Fawcett, Sreenivas Gollapudi, and Ali Kemal Sinop. 2021. Weighted Stackelberg Algorithms for Road Traffic Optimization. In *SIGSPATIAL '21: 29th International Conference on Advances in Geographic Information Systems, Virtual Event / Beijing, China, November 2-5, 2021*, Xiaofeng Meng, Fusheng Wang, Chang-Tien Lu, Yan Huang, Shashi Shekhar, and Xing Xie (Eds.). ACM, 57–68.

[33] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004), New York, NY, USA, May 17-20, 2004*, Jack S. Breese, Joan Feigenbaum, and Margo I. Seltzer (Eds.). ACM, 125–131.

[34] Patrice Marcotte. 1987. Algorithms for the Network Oligopoly Problem. *The Journal of the Operational Research Society* 38, 11 (1987), 1051–1065.

[35] Maren Martens and Martin Skutella. 2006. Flows on few paths: Algorithms and lower bounds. *Networks* 48, 2 (2006), 68–76.

[36] Carol A. Meyers and Andreas S. Schulz. 2012. The complexity of welfare maximization in congestion games. *Networks* 59, 2 (2012), 252–260. https://doi.org/10.1002/net.20439

[37] OpenStreetMap contributors. 2024. Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org.

[38] Ariel Orda, Raphael Rom, and Nahum Shimkin. 1993. Competitive Routing in Multi-User Communication Networks. In *Proceedings IEEE INFOCOM '93, The Conference on Computer Communications, Twelfth Annual Joint Conference of the IEEE Computer and Communications Societies, Networking: Foundation for the Future, San Francisco, CA, USA, March 28 - April 1, 1993.* IEEE Computer Society, 964–971.

[39] Arthur Cecil Pigou. 2013. *The Economics of Welfare.* Palgrave Macmillan.

[40] Ariel D. Procaccia. 2016. Cake Cutting Algorithms. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 311–330.

[41] Robert W. Rosenthal. 1973. The network equilibrium problem in integers. *Networks* 3 (1973), 53–59.

[42] Tim Roughgarden. 2002. How unfair is optimal routing?. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA*, David Eppstein (Ed.). ACM/SIAM, 203–204.

[43] Tim Roughgarden and Florian Schoppmann. 2015. Local smoothness and the price of anarchy in splittable congestion games. *J. Econ. Theory* 156 (2015), 317–342.

[44] Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259.

[45] Tim Roughgarden and Éva Tardos. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games Econ. Behav.* 47, 2 (2004), 389–403.

[46] Fernanda Salazar and Martin Skutella. 2009. Single-source k-splittable min-cost flows. *Oper. Res. Lett.* 37, 2 (2009), 71–74.