

Achieving Target Equilibria in Network Routing Games without Knowing the Latency Functions*

Umang Bhaskar[†] Katrina Ligett[‡] Leonard J. Schulman[§] Chaitanya Swamy[¶]

Abstract

The analysis of network routing games typically assumes precise and detailed information about the latency functions. Such information may, however, be unavailable or difficult to obtain. Moreover, one is often primarily interested in enforcing a desired target flow as the equilibrium by suitably influencing player behavior in the routing game. We ask whether one can achieve target flows as equilibria *without knowing the underlying latency functions*.

Our main result gives a crisp positive answer to this question. We show that, under fairly general settings, one can efficiently compute *edge tolls* that induce a given target multicommodity flow in a nonatomic routing game using a *polynomial number of queries* to an oracle. The oracle takes candidate tolls as input, and returns the resulting equilibrium flow. This result is obtained via a novel application of the ellipsoid method. Our algorithm extends easily to many other settings, such as (i) when certain edges cannot be tolled, or there is an upper bound on the total toll paid by a user, and (ii) general nonatomic congestion games. We obtain tighter bounds on the query complexity for series-parallel networks, and single-commodity routing games with linear latency functions, and complement these with a query-complexity lower bound. We also obtain strong positive results for *Stackelberg* routing to achieve target equilibria in series-parallel graphs.

Our results build upon various new techniques that we develop pertaining to the computation of, and connections between, different notions of approximate equilibrium; properties of multicommodity flows and tolls in series-parallel graphs; and sensitivity of equilibrium flow with respect to tolls. Our results demonstrate that one can indeed circumvent the potentially onerous task of modeling latency functions, and yet obtain meaningful results for the underlying routing game.

*An extended abstract of this work appeared in Proc. 55th FOCS 31-40, 2014 [8].

[†]School of Technology and Computer Science, Tata Institute of Fundamental Research, India. Work performed while at the Dept. of Computing and Mathematical Sciences, California Institute of Technology, and supported in part by a Linde/SISL post-doctoral fellowship and NSF grants CNS-0846025, CCF-1101470 and EPAS-1307794. Email: umang@tifrr.res.in.

[‡]Dept. of Computing and Mathematical Sciences, California Institute of Technology. Work supported in part by the Charles Lee Powell Foundation and a Microsoft Faculty Fellowship. Email: katrina@caltech.edu.

[§]Dept. of Computing and Mathematical Sciences, California Institute of Technology. Work supported in part by NSF grants 1038578, 1319745 and 1618795. Work performed in part at the Simons Institute for the Theory of Computing at UC Berkeley. Email: schulman@caltech.edu.

[¶]Dept. of Combinatorics and Optimization, University of Waterloo. Supported in part by NSERC grant 32760-06, an NSERC Discovery Accelerator Supplement Award, and an Ontario Early Researcher Award. Email: cswamy@math.uwaterloo.ca.

1 Introduction

Network routing games are a popular means of modeling settings where a collection of self-interested, uncoordinated users or agents route their traffic along an underlying network—prominent examples include communication and transportation networks—and have been extensively studied from various perspectives in the Transportation Science and Computer Science literature; see, e.g., [41, 4, 30, 31, 16, 23, 43, 44, 40], and the references therein. These games are typically described in terms of an underlying directed graph $G = (V, E)$, modeling the network; a set of commodities specified by source-sink pairs and the volume of traffic routed between them, modeling the different user-types; and latency functions or delay functions $(l_e^* : \mathbb{R}_+ \mapsto \mathbb{R}_+)_{e \in E}$ on the edges, with $l_e^*(x)$ modeling the delay experienced on edge e when volume x of traffic is present on it. The outcome of users’ strategic behavior is described by an *equilibrium* traffic pattern, where no user may unilaterally deviate and reduce her total delay.

The typical means of mathematically investigating network routing games takes the above specification as input, and thus, assumes, right at the onset, that one has precise, detailed information about the underlying latency functions. However, such precise information may be unavailable or hard to obtain, especially in large systems, without engaging in a highly non-trivial and potentially expensive modeling task. In fact, the task of capturing observed delays via suitable delay functions is a topic of much research in itself in fields such as queuing theory and transportation science. Recognizing that the modeling task of obtaining suitable latency functions is often really a means to facilitating a mathematical analysis of the underlying routing game, we ask whether one can sidestep this potentially demanding task and analyze the routing game *without knowing the underlying latency functions*. This is the question that motivates our work.

In routing games, there is often a central authority who has some limited ability to influence agents’ behavior by making suitable changes to the routing game, e.g., imposing tolls on the network edges. This influence can be used to alleviate the detrimental effects of selfish agent behavior, which might be expressed both in terms of the agents’ costs (i.e., price of anarchy) and externalities not captured by these (e.g., pollution costs in a road network). Thus, a natural and well-studied goal in network routing games is to *induce a desirable target traffic pattern as an equilibrium* by suitably influencing agents’ behavior. Such a target traffic pattern may be obtained by, e.g., limiting the traffic on every edge to a fraction of its capacity, or reducing the traffic near hospitals and schools. It is evident here that suitably modeling the latency functions is only a means to the end goal of achieving the target traffic pattern. Our work aims to shed light on the following question: *can one achieve this end without the means?* The only tool we allow ourselves in order to study the network, is the same intervention capability that will ultimately be used to create the preferred flow.

1.1 Our contributions

We initiate a systematic study of network routing games from the perspective of achieving target equilibria without knowing the latency functions. We introduce a *query model* for network routing games to study such questions, and obtain bounds on the query complexity of various tasks in this model.

The query model. We are explicitly given the underlying network $G = (V, E)$, the set of commodities specified by the source-sink pairs and the demands to be routed between them, and the *target multicommodity flow* f^* that we seek to achieve. We *do not*, however, know the underlying latency functions $(l_e^*)_{e \in E}$. Instead, the only information that we can glean about the latency functions is via queries to a *black box* or *oracle* (e.g., a simulation procedure) that outputs the equilibrium flow under a specified stimulus to the routing game. We investigate two methods for influencing agent behavior that have been considered extensively in the literature, which gives rise to two types of queries.

We primarily focus on the task of computing *edge tolls* to induce f^* (Sections 3 and 5.1). This yields the following query model: each query consists of a vector of tolls on the edges, and returns the equilibrium flow that results upon imposing these tolls. The goal is to minimize the number of queries required to compute tolls that yield f^* as the equilibrium.

We also explore, in Sections 4 and 5.2.2, the use of *Stackelberg routing* to induce f^* . Here, we control an α fraction of the total traffic volume. Each query is a Stackelberg routing, which is a flow of volume at most α times the total volume, and returns the equilibrium flow under this Stackelberg routing. The goal is to minimize the number of queries required to compute a Stackelberg routing that induces f^* as the equilibrium.

Our results and techniques. Our main result is a crisp and rather sweeping positive result showing that *one can always obtain tolls that induce a given target flow f^* with a polynomial number of queries* (Section 3.1). With linear latency functions, our algorithm computes tolls that enforce f^* exactly (Theorem 3.2). With more general latency functions, such as convex polynomial functions, equilibria may be irrational, so it is not meaningful to assume that a query returns the exact equilibrium. Instead, we assume that each query returns a (suitably-defined) approximate equilibrium and obtain tolls that enforce a flow that is component-wise close to f^* (Theorem 3.6).

The chief technical novelty underlying these results is an unconventional application of the ellipsoid method. We view the problem as one where we are searching for the (parameters of the) true latency functions l^* and tolls that induce f^* . It is information-theoretically *impossible*, however, to identify l^* (or even get close to it) in the query model since—as is the case even when G is a single edge—there may be no way of distinguishing two sets of latency functions. The key insight is that, notwithstanding this difficulty, if the current candidate tolls τ do not enforce f^* , then one can use the resulting equilibrium flow to identify a hyperplane that separates our current candidate (l, τ) from the true tuple (l^*, τ^*) . This enables one to use the machinery of the ellipsoid method to obtain tolls enforcing f^* in a polynomial number of queries.

Our ellipsoid-method based algorithm is quite versatile and can be easily adapted to handle various generalizations (Section 3.2). For instance, we can incorporate *any* linear constraints that tolls inducing f^* must satisfy, which one can separate over. This captures constraints where we disallow tolls on certain edges, or place an upper bound on the total toll paid by an agent. All our machinery extends seamlessly to the more-general setting of *nonatomic congestion games*. Finally, another notable extension is to the setting of *atomic splittable routing games* under the assumption that the equilibrium is unique.

In Sections 3.3 and 3.4, we devise algorithms with substantially improved query complexity for (a) multicommodity routing games on series-parallel (sepa) networks, and (b) single-commodity routing games on general networks, both with linear latency functions. For (a), we exploit the combinatorial structure of sepa graphs to design an algorithm with near-linear query complexity. We show that any toll-vector in a sepa graph can be converted into a simpler canonical form, which can be equivalently viewed in terms of certain labelings of the subgraphs of the sepa graph obtained via parallel joins; leveraging this yields an algorithm with near-linear query complexity. Our algorithm works more generally whenever we have an oracle that returns the (exact) equilibrium. For (b), we prove that (roughly speaking) the equilibrium flow is a linear function of tolls, and use linear algebra to infer the constants defining this linear map in $\tilde{O}(|E|^2)$ queries.

Complementing these upper bounds, we prove an $\Omega(|E|)$ *lower bound* (Theorem 5.1) on the query complexity of computing tolls that induce a target flow, even for single-commodity routing games on parallel-link graphs with linear delays. This almost matches the query complexity of our algorithm for sepa graphs.

En route to obtaining the above results, we prove various results that provide new insights into network routing games even in the standard non-black-box model where latency functions are known. For instance, we obtain results on: (a) the computation of approximate equilibria and their properties (Lemmas 3.4 and 3.5); (b) structural properties of tolls and multicommodity flows in sepa graphs (Section 3.3); and (c) sensitivity of equilibrium flow with respect to tolls (Theorem 3.24). We believe that these results and the machinery we develop to obtain them are of independent interest and likely to find various applications.

In Section 4, we investigate the use of Stackelberg routing to induce a given target flow. Stackelberg routing turns out to be significantly harder to leverage than edge tolls in the query model. This is perhaps not surprising given that designing effective Stackelberg routing strategies turns out to be more difficult than computing suitable edge tolls, even in the standard non-black-box setting where latency functions are given (see, e.g., [33, 9]). Nevertheless, we build upon the machinery that we develop for *sepa* graphs to give an efficient and general combinatorial algorithm that finds the desired Stackelberg routing using at most $|E|$ queries to an oracle returning equilibrium flows. This applies to any strictly increasing latency functions, and in particular, to linear latency functions. (Observe that this query complexity is even better than our query-complexity bound for inducing flows via tolls on *sepa* graphs.) Moreover, our algorithm determines the Stackelberg routing of smallest volume that can induce f^* .

We obtain various lower bounds in Section 5.2.2 that allude to the difficulty of computing a Stackelberg routing in general networks that induces a target flow. One possible strategy for finding such a Stackelberg routing is to use the queries to infer an (approximately) “equivalent” set of delay functions l , in the sense that any Stackelberg routing yields the same (or almost the same) resulting equilibrium under the two sets of delay functions. Then, since given the latency functions, it is easy to compute a Stackelberg routing that induces a target flow (see Lemma 2.2), one can find the desired Stackelberg routing. Theorem 5.8 shows that such an approach cannot work: in the query model, any algorithm that learns even an approximately equivalent set of delay functions must make an *exponential* number of queries. Theorem 5.10 proves an orthogonal computational lower bound showing that determining the equivalence of two given sets of latency functions is an *NP*-hard problem. As in the case of tolls, along the way, we uncover a new result about the hardness of Stackelberg routing. We show that the problem of finding a Stackelberg routing that minimizes the average delay of the remaining equilibrium flow is *NP*-hard to approximate within a factor better than $4/3$ (Theorem 5.13). The query complexity of finding a Stackelberg routing in general networks that induces a target flow remains an interesting open question for further research.

Our results on tolls and Stackelberg routing demonstrate that it is indeed possible to circumvent the potentially-onerous task of modeling latency functions, and yet obtain meaningful results for the underlying routing game. Our array of upper- and lower- bound results indicate the richness of the query model, and suggest a promising direction for further research.

1.2 Related work

Network routing/congestion games with nonatomic players—where each player controls an infinitesimal amount of traffic and there is a continuum of players—were first formally studied in the context of road traffic by Wardrop [41], and the equilibrium notion in such games is known as Wardrop equilibrium after him. Network routing games have since been widely studied in the fields of Transportation Science, Operations Research, and Computer Science; see, e.g., the monographs [30, 31] and the references therein. We limit ourselves to a survey of the results relevant to our work.

Equilibria are known to exist in network routing games, even with atomic players with splittable flow [28]. Nonatomic equilibria are known to be essentially unique, but this is not the case for atomic splittable routing games, where uniqueness criteria were recently obtained by Bhaskar et al. [6]. Equilibria in routing games are known to be inefficient, and considerable research in algorithmic game theory has focused on quantifying this inefficiency in terms of the *price of anarchy* (PoA) [24, 26] of the game, which measures, for a given objective, the worst-case ratio between the objective values of an equilibrium and the optimal solution. A celebrated result of Roughgarden [32], and Roughgarden and Tardos [37] gives tight bounds on the PoA for nonatomic routing games for the social welfare objective. Recently, similar results were obtained for the PoA in atomic splittable routing games [19, 36].

Given the inefficiency of equilibria, researchers have investigated ways of influencing player behavior so as to alleviate this inefficiency. The most common techniques studied to influence player behavior in network

congestion games are the imposition of tolls on the network edges, and Stackelberg routing. Network tolls are a classical means of congestion control, dating back to Pigou [27], and various results have demonstrated their effectiveness for both nonatomic routing [4, 10, 16, 23, 43] and atomic splittable routing [40, 44] showing that any minimal flow (in particular, an optimal flow) can be enforced via suitable efficiently-computable tolls. Stackelberg routing has also been well studied, and it is known that this is less effective in reducing the PoA. Whereas they can help in reducing the PoA to a constant for certain network topologies such as parallel-link graphs [33] and series-parallel graphs [40], it is known that this is not possible for general graphs [9]. Furthermore, it is known that it is *NP*-hard to compute the Stackelberg routing that minimizes the total cost at equilibrium, even for parallel-link graphs with linear delay functions [33]; a PTAS is known [25] for parallel-link graphs. All of these results pertain to the setting where one is given the latency functions.

To our knowledge, before our work, our query model was not studied in the literature. It is useful to contrast our query model with work in *empirical game theory*, which also studies games when players' costs are not explicitly given. In empirical game theory, each query specifies a (pure or mixed) strategy-profile, and returns the (expected) cost of each player under this strategy profile. In contrast, in our query model, we observe the equilibrium flow instead of individual player delays. This is more natural in the setting of routing games: in the absence of knowledge of the latency functions, one may only be able to calculate player delays under a strategy profile by routing players along the stipulated paths (and then observing player delays); but this may be infeasible since one cannot in fact impose routes on self-interested players. Moreover, whereas our goal is to obtain a desirable outcome as the equilibrium, the focus in empirical game theory is to compute an (approximate) equilibrium. Generic approaches to generate strategy-profiles for this purpose, and examples where these have proved useful are discussed by Wellman [42]. An oblivious algorithm that does not depend on player utilities, and instead uses best-responses to compute a pure Nash equilibrium in bimatrix games was given by Sureka and Wurman [39]. Starting with [35], various papers have studied the complexity of computing an exact or approximate correlated equilibrium in multi-player games using both pure- and mixed-strategy queries [2, 20, 21]. More recently, Fearnley et al. [14] study algorithms in the empirical-game-theory model for bimatrix games, congestion games, and graphical games, and obtain various bounds on the number of queries required for equilibrium computation.

A problem closely related to that studied in our paper is to efficiently learn parameters of the model from observing equilibrium behavior in response to stimuli. This is known as learning from revealed preferences. In the single-buyer problem, a single agent has a utility function over a set of goods, and chooses goods given a set of prices and a budget. Efficient algorithms for learning the utility function are known for this problem for various classes of utility functions, both in the statistical model when prices and budgets are drawn from a distribution [45, 3], and in the query model, when the algorithm can choose prices and budgets [3]. The problem of profit-maximization for a seller faced with such a buyer is studied by Amin et al. [1]. For linear utilities, they give an algorithm for which the regret of the seller decreases linearly in the number of queries. For the case of multiple buyers with linear utilities, when the response to prices for goods consists of aggregate demands from the buyers, Bei et al. give efficient algorithms that learn the utilities of the buyers [5]. Finally, Roth et al. [29] consider the problem of obtaining optimal outcomes in games, such as minimum-cost flow in nonatomic routing games and profit-maximizing prices in the single-buyer problem, when the parameters are unknown but equilibrium behavior is observed, and give efficient algorithms. Specialized to nonatomic routing games, their algorithm uses the inducement of a target flow as equilibrium as a subroutine. However, in contrast to our work, they use a tatonnement-like process to induce a given flow as equilibrium instead of the ellipsoid method.

2 Preliminaries and notation

A *nonatomic routing game* (or simply a routing game) is denoted by a tuple $\Gamma = (G, l, \mathcal{K})$, where $G = (V, E)$ is a directed graph with m edges and n nodes, $l = (l_e)_{e \in E}$ is a vector of latency or delay functions on the edges of the graph, and $\mathcal{K} = \{(s_i, t_i, d_i)\}_{i \leq k}$ is a set of k triples denoting sources, sinks, and demands for k commodities. The delay function $l_e : \mathbb{R}_+ \mapsto \mathbb{R}_+$ gives the delay on edge e as a function of the total flow on the edge. (Here, \mathbb{R}_+ is the set of nonnegative reals.) We assume that l_e is continuous and strictly increasing. For each commodity i , the demand d_i specifies the volume of flow that is routed from s_i to t_i by self-interested agents, each of whom controls an infinitesimal amount of flow and selects an s_i - t_i path as her strategy. The strategies selected by the agents thus induce a multicommodity flow $(f^i)_{i \leq k}$, where each $f^i = (f_e^i)_{e \in E}$ is an s_i - t_i flow of value d_i . That is, the vector $f^i = (f_e^i)_{e \in E}$ satisfies:

$$f^i \geq 0, \quad \sum_{(v,w) \in E} f_{vw}^i - \sum_{(u,v) \in E} f_{uv}^i = 0 \quad \forall v \in V \setminus \{s_i, t_i\}, \quad \sum_{(s,w) \in E} f_{sw}^i - \sum_{(u,s) \in E} f_{us}^i = d_i.$$

We call $f = (f^i)_{i \leq k}$ a feasible flow. We say that f is acyclic if $\{e : f_e^i > 0\}$ is acyclic for every commodity i . We overload notation and use f to also denote the total-flow vector $f = \sum_{i \leq k} f^i$. For a path P , we use $f_P > 0$ to denote $f_e > 0$ for all $e \in P$. We sometimes refer to $\bigcup_i \{s_i, t_i\}$ as the terminals of the routing game or multicommodity flow. Given an s - t flow f , we use $|f|$ to denote the value of f .

Let \mathcal{P}^i denote the collection of all s_i - t_i paths. Given a multicommodity flow $(f^i)_{i \leq k}$ induced by the agents' strategies, the delay of an agent that selects an s_i - t_i path P is the total delay, $l_P(f) := \sum_{e \in P} l_e(f_e)$, incurred on the edges of P . Each agent in a routing game seeks to minimize her own delay. To analyze the resulting strategic behavior, we focus on the concept of a *Nash equilibrium*, which is a profile of agents' strategies where no individual agent can reduce her delay by changing her strategy, assuming other agents do not change their strategies. In nonatomic routing games, this is formalized by the notion of *Wardrop equilibrium*.

Definition 2.1. A multicommodity flow \hat{f} is a *Wardrop equilibrium* (or simply an equilibrium) of a routing game Γ if it is feasible and for every commodity i , and all paths $P, Q \in \mathcal{P}^i$ with $\hat{f}_P^i > 0$, we have $l_P(\hat{f}) \leq l_Q(\hat{f})$.

A Wardrop equilibrium can be computed by solving the following convex program:

$$\min \Phi(f) := \sum_e \int_0^{f_e} l_e(x) dx \quad \text{s.t.} \quad f = \sum_{i=1}^k f^i, \quad f^i \text{ is an } s_i\text{-}t_i \text{ flow of value } d_i \quad \forall i = 1, \dots, k. \quad (1)$$

Given a routing game Γ and a feasible flow f , define $D^i(l, f) := \min_{P \in \mathcal{P}^i} l_P(f)$ for each commodity i , and call an edge e a shortest-path edge for commodity i with respect to f if e lies on some path $P \in \mathcal{P}^i$ such that $l_P(f) = D^i(l, f)$. Let $\mathcal{S}^i(l, f)$ be the set of shortest-path edges for commodity i with respect to f .

Tolls, Stackelberg routing, and our query model. We investigate both the use of edge tolls and Stackelberg routing to induce a given target flow. Tolls are additional costs on the edges that are paid by every player that uses the edge. A vector of tolls $\tau = (\tau_e)_e \in \mathbb{R}_+^E$ on the network edges thus changes the delay function on each edge e to $l_e^\tau(x) := l_e(x) + \tau_e$, and so the delay of an agent who chooses P is now $l_P(f) + \tau(P)$, where $\tau(P) := \sum_{e \in P} \tau_e$. We use $f(l, \tau)$ to denote the equilibrium flow obtained with delay functions $l = (l_e)_e$ and tolls $\tau = (\tau_e)_e$. We say that τ enforces a multicommodity flow f with latency functions l if the total flow $f(l, \tau)_e = f_e$ on every edge e .

For Stackelberg routing, in keeping with much of the literature, we focus on single-commodity routing games. Given a single-commodity routing game $\Gamma = (G, l, (s, t, d))$ and a parameter $\alpha \in [0, 1]$, a central

authority controls at most an α -fraction of the total s - t flow-volume d and routes this flow in any desired way, and then the remaining traffic routes itself selfishly. That is, a Stackelberg routing g is an s - t flow of value at most αd , which we call the Stackelberg demand. The Stackelberg routing g modifies the delay function on each edge e to $\tilde{l}_e(g; x) := l_e(x + g_e)$. The remaining $(1 - \alpha)d$ volume of traffic routes itself according to a Wardrop equilibrium, denoted by $f(l, g)$, of the instance $(G, \tilde{l}, (1 - \alpha)d)$. The total flow induced by a Stackelberg routing g is thus $g + f(l, g)$.

We shorten $f(l, \tau)$ to $f(\tau)$, and $f(l, g)$ to $f(g)$ when l is clear from the context.

In our query model, we are given the graph G , the commodity set $\mathcal{K} = \{(s_i, t_i, d_i)\}_{i \leq k}$, and a feasible target multicommodity flow f^* . There is an underlying routing game $\Gamma = (G, l^*, \mathcal{K})$, to which we are given query access. If our method of influencing equilibria is via tolls, then the oracle takes a toll-vector τ as input and returns the equilibrium flow $f(l^*, \tau)$ or a (suitably-defined) approximate equilibrium. Our goal is to minimize the number of queries required to compute tolls τ^* such that $f(l^*, \tau^*) = f^*$.

If our method of influencing equilibria is via Stackelberg routing, then we are also given the parameter $\alpha \in [0, 1]$. Each query takes a Stackelberg routing g with $|g| \leq \alpha d$ as input and returns the flow $f(l^*, g)$. Our goal is to minimize the number of queries required to compute a Stackelberg routing g^* of value at most αd such that $f(l^*, g^*) + g^* = f^*$, or determine that no such Stackelberg routing exists.

Properties of equilibria. The following facts about Wardrop equilibria, network tolls, and Stackelberg routing will be useful. Recall that the delay functions are nonnegative, continuous, and strictly increasing.

- A feasible flow f is an equilibrium flow iff $\sum_e (f_e - g_e) l_e(f_e) \leq 0$ for every feasible flow g ; see, e.g., [30]. Thus, the total-flow vector $(f_e)_e$ induced by an equilibrium flow is unique for strictly increasing delay functions.
- Every routing game admits an acyclic Wardrop equilibrium \hat{f} . If the delay functions are polynomial-time computable, then one can solve (1) and compute: (i) \hat{f} in polynomial time for linear delay functions; (ii) an acyclic flow f such that $\Phi(f) \leq \Phi(\hat{f}) + \epsilon$ in time $\text{poly}(\text{input size}, \log(\frac{1}{\epsilon}))$. See, e.g., [30], for details.
- Every minimal feasible flow f is enforceable via tolls [16, 23, 43], where f is minimal if there is no other feasible flow $g \neq f$ such that $g_e \leq f_e$ for every edge e . Given the edge delays $(l_e(f_e))_e$, these tolls can be computed by solving an LP, and are rational provided the commodity demands $(d_i)_i$ and the delays $(l_e(f_e))_e$ are rational.

The following lemma was essentially shown in [22]; we include a self-contained proof in Appendix A.

Lemma 2.2. *Let $(G, l, (d, s, t), \alpha)$ be a Stackelberg routing instance, and f^* be a feasible flow. Then, $f(g) + g = f^*$ for a Stackelberg routing g iff $g_e \leq f_e^*$ for every edge e , and $g_e = f_e^*$ for all $e \notin \mathcal{S}(l, f^*)$.*

Standard delay functions and encoding length. Our results hold for a broad class of underlying delay functions, that we now formally describe. We assume that we have an estimate U such that the target flow f^* , the parameters of the unknown true delay functions $(l_e^*)_e$, and the quantities that we seek to compute—tolls τ^* or the Stackelberg routing g^* inducing f^* —all have encoding length $O(\log U)$. So we may assume that every f_e^*, τ_e^*, g_e^* value is a multiple of $\frac{1}{U}$, and is at most U . Throughout, we use \mathcal{I} to denote the input size. We view U as part of the input, so $\mathcal{I} \geq \log U$ and is at least the encoding length of the given routing game.

When considering non-linear delay functions, we assume that the l_e^* s are convex functions with $\frac{dl_e^*(x)}{dx} \geq \frac{1}{U}$ for all $x \geq 0$. Although some of our results apply in more general settings (in particular, the ellipsoid-based method for computing tolls), we will often assume that the l_e^* s have a convenient parametric form, namely, they are (convex) polynomials of degree at most some known constant r . Given the $O(\log U)$ encoding length, we may assume that all coefficients lie in $[0, U]$ and are multiples of $\frac{1}{U}$. We refer to

such functions as *standard degree- r polynomials*. Under these conditions, it is easy to show (see Lemma 2.3) that there is some constant $K := K(r) = \text{poly}(U, \sum_i d_i)$ such that every delay function l_e^* satisfies

$$(x - y)(l_e^*(x) - l_e^*(y)) \leq \frac{\epsilon^2}{K} \implies |x - y| \leq \epsilon \quad \text{for all } x, y, \epsilon \geq 0 \quad (2)$$

$$|l_e^*(x) - l_e^*(y)| \leq K|x - y| \quad \text{for all } x, y \in [0, \sum_i d_i] \quad (3)$$

$$l_e^*(2x) \leq Kl_e^*(x) \quad \text{for all } x \geq 0 \quad (4)$$

These properties are referred to as *inverse- K -continuity*, *K -Lipschitz*, and *K -growth-boundedness* respectively.

Lemma 2.3. *Let $l(x) = a_0 + a_1x + \dots + a_rx^r$ be a convex degree- r polynomial such that $a_1 > 0$, and all a_i s lie in $[0, U]$ and are multiples of $\frac{1}{U}$. Then l satisfies (2)–(4) with $K = \max\{U, 2^r, rU(\sum_i d_i)^{r-1}\}$.*

Proof. Let $l'(x) := \frac{dl(x)}{dx}$ denote the derivative of l . Since l is convex, we have $|l(x) - l(y)| \geq |x - y| \cdot l'(\min\{x, y\}) \geq |x - y| \cdot l'(0) \geq |x - y|/U$. Therefore, $\frac{(x-y)^2}{U} \leq (x - y)(l(x) - l(y)) \leq \frac{\epsilon^2}{K}$ and so $|x - y| \leq \epsilon$.

Again, by convexity, $|l(x) - l(y)| \leq |x - y| \cdot l'(\max\{x, y\})$ and $l'(z) \leq rU(\sum_i d_i)^{r-1} \leq K$ for all $z \leq \sum_i d_i$.

Finally, it is clear that $l(2x) \leq 2^r l(x) \leq Kl(x)$ for all $x \geq 0$. ■

Asymptotic notation, NP-hard-ness, and inapproximability Our goal in this paper is to obtain asymptotic bounds, both upper and lower, on the problem of enforcing a target flow. While our primary concern is the number of queries, in Section 5 we will also be concerned with computational complexity. We frequently employ asymptotic notation to denote the complexity (query or time) of our algorithms and lower bounds. In particular, the notation $\tilde{O}(f(n))$ is used to denote a function that asymptotically grows no faster than $cf(n) \log^k n$ for some constants c, k , independent of the input. That is, the polylogarithmic factors are “hidden”. A good reference for a fuller introduction to asymptotic notation, and other concepts in this subsection, is [12].

Unlike query complexity, where unconditional lower bounds can be obtained frequently, for computational complexity very few unconditional lower bounds are known. Instead, a standard practice for showing computational lower bounds for a problem is to show that the problem is *NP-hard*. The class of *NP-hard* problems consists of those for which no polynomial-time algorithms are known, and in fact a polynomial-time algorithm for a single such problem would give polynomial-time algorithms for the entire related class of NP-complete problems. The standard method for showing a problem (say P) is *NP-hard* is to give a polynomial-time reduction from an existing *NP-hard* problem (say Q). Thus if one can solve P in polynomial time, one can also solve Q, which was known to be *NP-hard*. We use this in Section 5. For *NP-hard* problems, since it is unlikely that an exact solution can be obtained in polynomial time, one can then ask if an approximate solution with a particular approximation ratio — with approximation ratio appropriately defined — can be obtained in polynomial time. We can extend the notion of *NP-hard-ness* to hardness of obtaining such approximate solutions as well. The class of *APX-hard* problems are those for which there exists a constant, such that it is *NP-hard* to obtain a solution with an approximation ratio better than the constant.

3 Inducing target flows via tolls

Recall that here we seek to compute tolls that enforce a given target flow f^* given black-box access to a routing game $\Gamma^* = (G, l^*, (s_i, t_i, d_i)_{i \leq k})$, i.e., without knowing l^* . Our main result is a crisp positive result showing that we can always achieve this end with a polynomial number of queries by leveraging the

ellipsoid method in a novel fashion (Section 3.1). Our algorithm computes tolls that enforce: (a) f^* exactly, for standard linear latency functions (where it is reasonable to assume that the black box returns the exact equilibrium); and (b) a flow that is component-wise close to f^* , for standard polynomial functions, where we now assume that each query only returns an approximate equilibrium (see Definition 3.3). The main idea here is to view the parameters of the latency functions and the tolls as variables, and use the ellipsoid method to search for the tuple (l^*, τ^*) , where τ^* is such that $f(l^*, \tau^*) = f^*$. The key observation is that although we cannot hope to nail down l^* , given a candidate (l, τ) such that $f(l, \tau) \neq f^*$, one can derive a hyperplane separating (l, τ) from (l^*, τ^*) using f^* and the equilibrium flow $f(l, \tau)$ returned by our oracle.

We showcase the versatility of our algorithm by showing that it is easily adapted to handle various extensions (Section 3.2). For instance, we can impose *any* linear constraints on tolls given by a separation oracle; examples include the constraint that certain edges cannot be tolled or that the total toll paid by a user is at most a given budget. Other notable extensions include the extension to general nonatomic congestion games, and to atomic splittable routing games under the assumption that the equilibrium is unique.

Finally, we devise algorithms with significantly improved query complexity for multicommodity routing games on series-parallel (sepa) networks (Section 3.3), and single-commodity routing games on general networks (Section 3.4), both with linear latency functions. We exploit the combinatorial structure of sepa graphs to design an algorithm with near-linear query complexity, which almost matches the linear lower bound shown in Theorem 5.1 for even parallel-link graphs with linear latencies. For single-commodity routing games on general graphs with linear latencies, we show that flows are linear functions of tolls and infer this linear map using $\tilde{O}(m^2)$ queries.

3.1 An ellipsoid-method based algorithm for general routing games

The ellipsoid method for finding a feasible point starts by containing the feasible region within a ball and generates a sequence of ellipsoids of successively smaller volumes. In each iteration, one examines the center of the current ellipsoid. If this is infeasible, then one uses a violated inequality to obtain a hyperplane, called a separating hyperplane, to separate the current ellipsoid center from the feasible region. One then generates a new ellipsoid by finding the minimum-volume ellipsoid containing the half of the current ellipsoid that includes the feasible region. We utilize the following well-known theorem about the ellipsoid method.

Theorem 3.1 ([18]). *Let $X \subseteq \mathbb{R}^n$ be a polytope described by constraints having encoding length at most M . Suppose that for each $y \in \mathbb{R}^n$, we can determine if $y \notin X$ and if so, return a hyperplane of encoding length at most M separating y from X . Then, we can use the ellipsoid method to find a point $x \in X$ or determine that $X = \emptyset$ in time $\text{poly}(n, M)$.*

Linear latencies. We first consider the case where each latency function $l_e^*(x)$ is a standard linear function of the form $a_e^*x + b_e^*$, and our black box returns the exact equilibrium flow induced by the input (rational) tolls. Thus, for every e , $a_e^* \in (0, U)$, $b_e^* \in [0, U]$, and a_e^*, b_e^* are multiples of $\frac{1}{U}$. In a somewhat atypical use of the ellipsoid method, we use the ellipsoid method to search for the point $(a_e^*, b_e^*, \tau_e^*)_e$. Abusing notation slightly, for a linear latency function $l(x) = ax + b$, we use l to also denote the tuple (a, b) .

Theorem 3.2. *Given a target acyclic multicommodity flow f^* and query access to Γ^* , we can compute tolls that enforce f^* or determine that no such tolls exist, in polynomial time using a polynomial number of queries.*

Proof. We utilize the ellipsoid method and Theorem 3.1. Given the center $(\hat{l} = (\hat{a}_e, \hat{b}_e)_e, \hat{\tau})$ of the current ellipsoid, we first check if $\hat{a}, \hat{b}, \hat{\tau} \geq 0$, and if not, use the violated constraint as the separating hyperplane. Next, we use the black box to obtain $g = f(l^*, \hat{\tau})$. If $g = f^*$, then we are done. Otherwise, we obtain

a separating hyperplane of encoding length $\text{poly}(\mathcal{I})$ as follows. (Note that the encoding length of $(\hat{l}, \hat{\tau})$ is $\text{poly}(\mathcal{I})$.) We consider two cases.

Case 1: $f(\hat{l}, \hat{\tau}) \neq f^*$. Note that we can determine this without having to compute the equilibrium flow $f(\hat{l}, \hat{\tau})$. Since f^* is acyclic, we can efficiently find a commodity i , and s_i - t_i paths P, Q such that $f_P^* > 0$ and $\hat{l}_P(f^*) + \hat{\tau}(P) > \hat{l}_Q(f^*) + \hat{\tau}(Q)$. But since $f^* = f(l^*, \tau^*)$, we also have $l_P^*(f^*) + \tau^*(P) \leq l_Q^*(f^*) + \tau^*(Q)$. Thus, the inequality

$$l_P(f^*) + \tau(P) \leq l_Q(f^*) + \tau(Q)$$

where the parameters of l and τ are variables yields the desired separating hyperplane.

Case 2: $f(\hat{l}, \hat{\tau}) = f^*$. Now since $g \neq f^*$ and is acyclic, we can again find efficiently a commodity i and paths $P, Q \in \mathcal{P}^i$ such that $g_P > 0$ and $\hat{l}_P(g) + \hat{\tau}(P) > \hat{l}_Q(g) + \hat{\tau}(Q)$. Since $g = f(l^*, \hat{\tau})$, we also have $l_P^*(g) + \hat{\tau}(P) \leq l_Q^*(g) + \hat{\tau}(Q)$. Thus, the inequality $l_P(g) + \hat{\tau}(P) \leq l_Q(g) + \hat{\tau}(Q)$, where now *only* the l_e s are variables, yields the desired separating hyperplane. ■

Polynomial latency functions and approximate equilibria. We now consider the setting where the latency functions $(l_e^*)_e$ are standard degree- r polynomials, where r is a known constant. As before, we also use l to denote the tuple of coefficients of the polynomial given by l . Since the Wardrop equilibrium may now require irrational numbers, it is unreasonable to assume that a query returns the equilibrium flow. So we assume that our black box returns an acyclic approximate equilibrium and show that we can nevertheless compute tolls that induce an equilibrium that is component-wise close to f^* . We first define approximate equilibria. Recall that $D^i(l, f) = \min_{P \in \mathcal{P}^i} l_P(f)$, and given tolls τ , we define $l_e^\tau(x) := l_e(x) + \tau_e$.

Definition 3.3. We say that a feasible flow f is an ϵ -approximate equilibrium, or simply an ϵ -equilibrium, of a routing game $(G, l, (s_i, t_i, d_i)_{i \leq k})$ if $\sum_e f_e l_e(f_e) \leq \sum_i d_i (D^i(l, f) + \epsilon)$.

Notice that our approximate-equilibrium notion is implied by the more-stringent (and oft-cited) condition requiring that if $f_P > 0$ for $P \in \mathcal{P}^i$ then $l_P(f) \leq D^i(l, f) + \epsilon$. Importantly, our notion turns out to be weak enough that one can argue that an acyclic ϵ -equilibrium can be computed in time $\text{poly}(\mathcal{I}, \log(\frac{1}{\epsilon}))$ for any $\epsilon > 0$, which lends credence to our assumption that the black box returns an acyclic ϵ -equilibrium, and yet is strong enough that one can leverage it within the framework of the ellipsoid method (see Theorem 3.6). Unless otherwise stated, when we refer to a routing game below, we assume that the latency functions satisfy the mild conditions (2)–(4), with $\log K$ being polynomially bounded. The following Lemma is proved in Appendix B.

Lemma 3.4. Given a routing game with polynomial-time computable latency functions, one can compute an acyclic ϵ -equilibrium in time $\text{poly}(\mathcal{I}, \log(\frac{1}{\epsilon}))$.

Lemma 3.5. Let \hat{f} be a Wardrop equilibrium and g be an ϵ -equilibrium of a routing game $(G, l, (s_i, t_i, d_i)_{i \leq k})$. Then, $\|g - \hat{f}\|_\infty := \max_e |g_e - \hat{f}_e| \leq \sqrt{K\epsilon \sum_i d_i}$.

Proof. We have $\sum_e g_e l_e(g_e) \leq \sum_i d_i (D^i(l, g) + \epsilon)$ and $\sum_e \hat{f}_e l_e(g_e) \geq \sum_i d_i D^i(l, g)$. So $\sum_e (g_e - \hat{f}_e) l_e(g_e) \leq \epsilon \sum_i d_i$. Also, $\sum_e (\hat{f}_e - g_e) l_e(\hat{f}_e) \leq 0$. So $\sum_e (g_e - \hat{f}_e) (l_e(g_e) - l_e(\hat{f}_e)) \leq \epsilon \sum_i d_i$. Each term of this summation is nonnegative and hence, at most $\epsilon \sum_i d_i$; therefore, $|g_e - \hat{f}_e| \leq \sqrt{K\epsilon \sum_i d_i}$ by inverse- K -continuity. ■

Define an ϵ -oracle for tolls to be an oracle that receives tolls $\tau \in \mathbb{R}_+^E$ as input and returns an ϵ -equilibrium of the routing game $(G, l^{*\tau}, (s_i, t_i, d_i)_{i \leq k})$ having encoding length $\text{poly}(\mathcal{I}, \log(\frac{1}{\epsilon}))$.

Theorem 3.6. Let f^* be a target acyclic multicommodity flow and $\delta > 0$. Let $\epsilon = \frac{\delta^2}{Kmk \sum_i d_i}$. Then, in time $\text{poly}(\mathcal{I}, \log(\frac{1}{\delta}))$ and using $\text{poly}(\mathcal{I}, \log(\frac{1}{\delta}))$ ϵ -oracle queries, we can compute tolls τ such that $\|f(l^*, \tau) - f^*\|_\infty \leq 2\delta$ or determine that no such tolls exist.

Proof. As before, we use the ellipsoid method. Let $(\hat{l}, \hat{\tau})$ be the center of the current ellipsoid. Assume that $\hat{l}, \hat{\tau} \geq 0$ and each function \hat{l}_e has slope at least $\frac{1}{U}$; otherwise, we can use a violated constraint as the separating hyperplane. We use the oracle with toll-vector $\hat{\tau}$ to obtain an acyclic ϵ -equilibrium flow g . Then, we have $\|g - f(l^*, \hat{\tau})\|_\infty \leq \sqrt{K\epsilon \sum_i d_i} = \delta/\sqrt{mk}$ by Lemma 3.5.

We can efficiently determine if $f(\hat{l}, \hat{\tau}) \neq f^*$, and if so, then as in Case 1 in the proof of Theorem 3.2, we can obtain a separating hyperplane of encoding length $\text{poly}(\mathcal{I})$. So assume otherwise.

Now we check if g is an $mk\epsilon$ -equilibrium for the latency functions $(\hat{l}_e^{\hat{\tau}})_e$. If so, then $\|g - f^*\|_\infty \leq \delta$ and so $\|f(l^*, \hat{\tau}) - f^*\|_\infty \leq 2\delta$ and we are done. Otherwise, we find a valid path-decomposition $x = (x_{i,P})_{i,P \in \mathcal{P}^i}$ of g having support of size at most mk . That is, we have $x \geq 0$, $\sum_{P \in \mathcal{P}^i} x_{i,P} = d_i$ for every commodity i , $\sum_i \sum_{P \in \mathcal{P}^i: e \in P} x_{i,P} = g_e$ for all e , and $\sum_i |\{P : x_{i,P} > 0\}| \leq mk$. We may assume that every non-zero $x_{i,P}$ value has encoding length that is polynomial in \mathcal{I} and the size of g . Then

$$\sum_i \sum_{P \in \mathcal{P}^i} x_{i,P} (\hat{l}_P^{\hat{\tau}}(g) - D^i(\hat{l}^{\hat{\tau}}, g)) = \sum_e g_e \hat{l}_e^{\hat{\tau}}(g_e) - \sum_i d_i D^i(\hat{l}^{\hat{\tau}}, g) > mk\epsilon \sum_i d_i$$

where the last inequality follows since g is not an $mk\epsilon$ -equilibrium for $(\hat{l}_e^{\hat{\tau}})_e$. Since the support of x has size at most mk , this implies that there is some commodity j and some path $R \in \mathcal{P}^j$ such that $x_{j,R}(\hat{l}_R^{\hat{\tau}}(g) - D^j(\hat{l}^{\hat{\tau}}, g)) > \epsilon \sum_i d_i$. Moreover, we can find such a j and path $R \in \mathcal{P}^j$ efficiently by simply enumerating the paths in the support of x . Let $Q \in \mathcal{P}^j$ be such that $\hat{l}_Q^{\hat{\tau}}(g) = D^j(\hat{l}^{\hat{\tau}}, g)$.

Since g is an ϵ -equilibrium for the latency functions $(l_e^{\hat{\tau}})_e$, again considering the path-decomposition x , we have $\sum_i \sum_{P \in \mathcal{P}^i} x_{i,P} (l_P^{\hat{\tau}}(g) - D^i(l^{\hat{\tau}}, g)) \leq \epsilon \sum_i d_i$. Each term in this sum is nonnegative, so each term is at most $\epsilon \sum_i d_i$. In particular, we have $x_{j,R} (l_R^{\hat{\tau}}(g) - l_Q^{\hat{\tau}}(g)) \leq x_{j,R} (l_R^{\hat{\tau}}(g) - D^j(l^{\hat{\tau}}, g)) \leq \epsilon \sum_i d_i$. So the inequality $x_{j,R} (l_R(g) + \hat{\tau}(R) - l_Q(g) - \hat{\tau}(Q)) \leq \epsilon \sum_i d_i$, with l_e s as the variables, is valid for (l^*, τ^*) but is violated by $(\hat{l}, \hat{\tau})$. This yields a separating hyperplane of encoding length $\text{poly}(\mathcal{I}, \log(\frac{1}{\epsilon}))$. ■

General latency functions. We now consider the setting where the l_e^* s are arbitrary nonnegative, continuous, convex, strictly increasing functions. Let $D = \sum_i d_i$. Abstracting away the properties underlying standard polynomial functions, we assume that we have estimates $U \in \mathbb{Z}_+$, and $\sigma_\gamma \in (0, 1]$ for all $\gamma > 0$ such that the l_e^* s satisfy:

$$l_e^*(x) \in [0, U], \quad \frac{dl_e^*(x)}{dx} \in \left[\frac{1}{U}, U \right], \quad l_e^*((1+\sigma_\gamma)x) \leq (1+\gamma)l_e^*(x) \quad \forall \gamma > 0 \quad \text{for all } x \in [0, D] \quad (5)$$

It is not hard to see that the bounds on $\{\frac{dl_e^*}{dx}\}_e$ imply that the l_e^* s satisfy (2), (3) with $K = U$; the third condition above can be viewed as a refinement of (4). We assume (as before) that every f_e^* and τ_e^* value is a multiple of $\frac{1}{U}$; we also assume that D is a multiple of $\frac{1}{U}$. As before, we assume that U is specified with the input so $\mathcal{I} \geq \log U$.

Let $\delta > 0$ be a parameter specifying that the tolls τ we compute should satisfy $\|f(l^*, \tau) - f^*\|_\infty = O(\delta)$. We will adapt our ellipsoid-based algorithm to obtain an algorithm with query-complexity $\text{poly}(\mathcal{I}, \frac{1}{\sigma_\gamma}, \log(\frac{1}{\delta}))$, where $\gamma = \text{poly}(\frac{\delta}{UDmk})$. Unlike before, we do not now have a convenient parametric representation for the l_e^* s, so the chief new ingredient here is to come up with a compact discretized representation of the l_e^* s that does not incur much loss, and then use the ellipsoid method to infer this representation (or obtain suitable tolls enroute). We first describe the discretized representation (Definitions 3.7, 3.8) and show that it satisfies some useful properties (Lemma 3.9), and then show how we can leverage the ellipsoid method to work with this representation (Theorem 3.11).

For any $B > 0$ and $x \in \mathbb{R}$, we use $\lfloor x \rfloor_B := \lfloor Bx \rfloor / B$ to denote the largest multiple of $1/B$ that is at most x . Similarly, $\lceil x \rceil_B := \lceil Bx \rceil / B$ denotes the smallest multiple of $1/B$ that is at least x . Set $\epsilon = \frac{\delta^2}{UDmk}$

and $\gamma = \frac{\epsilon}{UDm}$, and let $\sigma = \sigma_\gamma$. Let $U_1 = U \lceil m/\epsilon \rceil$ and $U_2 = U_1 U$. Roughly speaking, the discretized version of l_e^* records $\lceil l_e^*(x) \rceil_{U_2}$ for all points of the form $\frac{(1+\sigma)^i}{U_1}$ in $[0, D]$. More precisely, define

$$X = \{0, D\} \cup \{f_e^* : e \in E\} \cup \left\{ \left\lfloor \frac{(1+\sigma)^i}{U_1} \right\rfloor_{U_1}, \left\lceil \frac{(1+\sigma)^i}{U_1} \right\rceil_{U_1} : i \geq 0, \frac{(1+\sigma)^i}{U_1} \leq D \right\}.$$

Observe that $M := |X| = O\left(\frac{\log(DU_1)}{\sigma}\right) = \text{poly}\left(\mathcal{I}, \frac{1}{\sigma}\right)$, all points in X are multiples of $\frac{1}{U_1}$, and $X \subseteq [0, D]$ since D is a multiple of $\frac{1}{U_1}$. Let $X = \{x_1, x_2, \dots, x_M\}$, where $x_1 = 0 < x_2 < \dots < x_M = D$. Note that $x_2 \leq x_1 + \frac{1}{U_1}$ and $x_{i+1} \leq (1 + \sigma)x_i$ for all $i = 2, \dots, M - 1$.

Definition 3.7. Let $l : \mathbb{R}_+ \mapsto \mathbb{R}_+$. The *discretization* of l is specified by the vector $l^\perp := (\lceil l(x) \rceil_{U_2})_{x \in X}$.

Definition 3.8. Given any $v \in \mathbb{R}_+^X$, we define the *interpolation* of v , denoted by v^\vee , to be the following piecewise-linear function with domain $[0, D]$:

$$v^\vee(x) := v_{x_i} + \frac{v_{x_{i+1}} - v_{x_i}}{x_{i+1} - x_i} \cdot (x - x_i) \quad \forall x \in [x_i, x_{i+1}], \forall i = 1, \dots, M - 1.$$

Lemma 3.9. Let $l : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be a function satisfying (5). Let $l^\perp \in \mathbb{R}_+^X$ be the discretization of l , and $\tilde{l} : [0, D] \mapsto \mathbb{R}_+$ be the interpolation of l^\perp (i.e., $\tilde{l} = v^\vee$). Then, the following hold.

- (i) $l(x) \leq \tilde{l}(x) \leq U$ for all $x \in [0, D]$
- (ii) $\tilde{l}(x) \leq l(x) + \frac{1}{U_2}$ for all $x \in X$, and $\tilde{l}(x) \leq (1 + \gamma)l(x) + \frac{U}{U_1} + \frac{1}{U_2}$ for all $x \in [0, D] \setminus X$
- (iii) The slope of each linear segment of \tilde{l} lies in $[\frac{1}{U}, U]$.

Proof. We have $l(x) \leq l_x^\perp \leq l(x) + \frac{1}{U_2}$ for all $x \in X$ by definition. Also, since $l(x) \leq U$ and U is a multiple of $\frac{1}{U_2}$, we have $l_x^\perp \leq U$ for all $x \in X$. For every $x \in [0, D]$, $\tilde{l}(x)$ is either l_x^\perp or a convex combination $\alpha l_a^\perp + (1 - \alpha)l_b^\perp$ for some $\alpha \in [0, 1]$, $a, b \in X$ such that $x = \alpha a + (1 - \alpha)b$. In the former case, we have already shown that $l(x) \leq \tilde{l}(x) \leq U$. In the latter case, we again have $\tilde{l}(x) \leq U$, and since l is a convex, we have $\tilde{l}(x) \geq \alpha l(a) + (1 - \alpha)l(b) \geq l(x)$. This proves part (i).

For part (ii), for $x \in X$, we have $\tilde{l}(x) = l_x^\perp \leq l(x) + \frac{1}{U_2}$. For $x \in [0, D] \setminus X$, suppose $x \in (x_i, x_{i+1})$. Then

$$\tilde{l}(x) \leq \max\left\{l_{x_i}^\perp, l_{x_{i+1}}^\perp\right\} \leq l(x_{i+1}) + \frac{1}{U_2} \leq l\left((1+\sigma)x_i + \frac{1}{U_1}\right) + \frac{1}{U_2} \leq (1+\gamma)l(x_i) + \frac{U}{U_1} + \frac{1}{U_2} \leq (1+\gamma)l(x) + \frac{1}{U_2}$$

where the penultimate inequality follows from (5).

For part (iii), consider a linear segment of \tilde{l} between $a = x_i, b = x_{i+1} \in X$. We have $U_2 \tilde{l}(b) \geq U_2 l(b) \geq U_2 l(a) + U_2(b - a)/U$ since $\frac{dl(x)}{dx} \geq \frac{1}{U}$ for all $x \in [0, D]$. Now $U_2 \tilde{l}(b)$ and $U_2(b - a)/U$ are integers, so we also have $U_2 \tilde{l}(b) \geq \lceil U_2 l(a) \rceil + U_2(b - a)/U$, or equivalently that $\tilde{l}(b) - \tilde{l}(a) \geq (b - a)/U$. Similarly, we have $U_2 l(b) \leq U_2 \tilde{l}(a) + U_2 U(b - a)$. Therefore, $\lceil U_2 l(b) \rceil \leq U_2 \tilde{l}(a) + U_2 U(b - a)$, or equivalently, $\tilde{l}(b) - \tilde{l}(a) \leq U(b - a)$. ■

Our goal now is to search for the point $(l^{*\perp}, \tau^*) := (l_e^{*\perp}, \tau_e^*)_e$. Note that the encoding length of $(l^{*\perp}, \tau^*)$ is $\text{poly}\left(\mathcal{I}, \frac{1}{\sigma}, \log\left(\frac{1}{\delta}\right)\right)$. Let \tilde{l}_e^* denote the interpolation of $l_e^{*\perp}$ for every edge e . Slightly abusing notation, we say that a piecewise linear function l satisfies $\frac{dl_e(x)}{dx} \in [\frac{1}{U}, U]$ if the slope of every linear segment of l lies in $[\frac{1}{U}, U]$. We begin by observing that Lemma 3.5 continues to hold for latency functions satisfying (5), since we only utilize (2) (i.e., inverse K -continuity) for $x, y \in [0, D]$ in the proof, which holds with $K = U$ if $\frac{dl_e(x)}{dx} \geq \frac{1}{U}$ for all $x \in [0, D]$.

Claim 3.10. Let \hat{f} be a Wardrop equilibrium and g be an ϵ -equilibrium of a routing game $(G, l, (s_i, t_i, d_i)_{i \leq k})$, where the l_e s satisfy $\frac{dl_e(x)}{dx} \geq \frac{1}{U}$ for all $x \in [0, \sum_i d_i]$. Then, $\|g - \hat{f}\|_\infty := \max_e |g_e - \hat{f}_e| \leq \sqrt{U\epsilon \sum_i d_i}$.

Theorem 3.11. Let f^* be a target acyclic multicommodity flow f^* and $\delta > 0$. Let $\epsilon = \frac{\delta^2}{UDmk}$, $\gamma = \frac{\epsilon}{UDm}$ and $\sigma = \sigma_\gamma$. Then, in time $\text{poly}(\mathcal{I}, \frac{1}{\sigma}, \log(\frac{1}{\delta}))$ and using $\text{poly}(\mathcal{I}, \frac{1}{\sigma}, \log(\frac{1}{\delta}))$ ϵ -oracle queries, we can compute tolls τ such that $\|f(l^*, \tau) - f^*\|_\infty \leq 5\delta$ or determine that no such tolls exist.

Proof. Again, we use the ellipsoid method and dovetail the proof of Theorem 3.6. Let $(\hat{l}, \hat{\tau})$ be the center of the current ellipsoid. Let $\tilde{l} = \iota^{\hat{l}}$ denote the interpolation of \hat{l} . We first check that $\hat{\tau} \geq 0$, and for every e , the following hold.

$$\hat{l}_{e,x} \in [0, U] \quad \forall x \in X, \quad \frac{x_{i+1} - x_i}{U} \leq \hat{l}_{e,x_{i+1}} - \hat{l}_{e,x_i} \leq U(x_{i+1} - x_i) \quad \forall i = 1, \dots, M-1.$$

Parts (i) and (iii) of Lemma 3.9 imply that $(l^{*\perp}, \tau^*)$ satisfies the above constraints. So if any of the above constraints do not hold, then we can use the violated constraint as a separating hyperplane. So assume otherwise. Then $\tilde{l}(x) \in [0, U]$ and $\frac{d\tilde{l}(x)}{dx} \in [\frac{1}{U}, U]$ for all $x \in [0, D]$. (But \tilde{l} need not be convex.) Note that the equilibrium flow $f(\tilde{l}, \hat{\tau})$ is well defined.

We determine if there is a commodity i and paths $P, Q \in \mathcal{P}^i$ such that $f_P^* > 0$ and $\tilde{l}_P(f^*) + \hat{\tau}(P) > \tilde{l}_Q(f^*) + \hat{\tau}(Q) + \epsilon$. (As before, we can do this efficiently since f^* is acyclic.) Since $f^* = f(l^*, \tau^*)$, we also have $l_P^*(f^*) + \tau^*(P) \leq l_Q^*(f^*) + \tau^*(Q)$, which implies that $\tilde{l}_P^*(f^*) + \tau^*(P) \leq \tilde{l}_Q^*(f^*) + \tau^*(Q) + \epsilon$ by part (ii) of Lemma 3.9 since $f_e^* \in X$ for all e and $\frac{1}{U_2} \leq \frac{\epsilon}{m}$. Since $f_e^* \in X$ for all e , we have $\tilde{l}_e(f_e^*) = \hat{l}_{e,f_e^*}$ and $\tilde{l}_e^*(f_e^*) = l_{e,f_e^*}^*$. Thus, in this case, the inequality

$$\sum_{e \in P} (l_{e,f_e^*} + \tau_e) \leq \sum_{e \in Q} (l_{e,f_e^*} + \tau_e) + \epsilon$$

where the l_{e,f_e^*} s and the τ_e s are variables yields the desired separating hyperplane of encoding length $\text{poly}(\mathcal{I}, \frac{1}{\sigma}, \log(\frac{1}{\delta}))$.

So assume that the above case does not happen. Then f^* is an ϵ -equilibrium for the latency functions $(\tilde{l}_e^{\hat{\tau}})_e$, and so $\|f(\tilde{l}, \hat{\tau}) - f^*\|_\infty \leq \sqrt{UD\epsilon} = \delta/\sqrt{mk}$ by Claim 3.10. We use the oracle with toll-vector $\hat{\tau}$ to obtain an acyclic ϵ -equilibrium flow g . Then, we have $\|g - f(l^*, \hat{\tau})\|_\infty \leq \sqrt{UD\epsilon} = \delta/\sqrt{mk}$ by Claim 3.10. Now we check if g is an $6mk\epsilon$ -equilibrium for the latency functions $(\tilde{l}_e^{\hat{\tau}})_e$. If so, then

$$\|g - f^*\|_\infty \leq \|g - f(\tilde{l}, \hat{\tau})\|_\infty + \|f(\tilde{l}, \hat{\tau}) - f^*\|_\infty \leq \sqrt{6UDmk\epsilon} + \sqrt{UD\epsilon} \leq 4\delta$$

and so $\|f(l^*, \hat{\tau}) - f^*\|_\infty \leq 5\delta$ and we are done.

Otherwise, as in the proof of Theorem 3.6, if $x = (x_{i,P})_{i,P \in \mathcal{P}^i}$ is a path-decomposition of g having support of size at most mk , then we can find some commodity j and paths $R, Q \in \mathcal{P}^j$ such that $x_{j,R}(\tilde{l}_R^{\hat{\tau}}(g) - \tilde{l}_Q^{\hat{\tau}}(g)) > 6\epsilon D$. Also, as before, since g is an ϵ -equilibrium for the latency functions $(l_e^{\hat{\tau}})_e$, we have $x_{j,R}(l_R^{\hat{\tau}}(g) - l_Q^{\hat{\tau}}(g)) \leq \epsilon D$. By part (ii) of Lemma 3.9 and since $\frac{U}{U_1}, \frac{1}{U_2} \leq \frac{\epsilon}{m}$ and $x_{j,R} \leq D$, we have $x_{j,R} \sum_{e \in R} (\tilde{l}_e^{\hat{\tau}}(g_e) + \hat{\tau}_e) \leq (1 + \gamma)x_{j,R} l_R^{\hat{\tau}}(g) + 2\epsilon D$. Since we also have $\tilde{l}_e^{\hat{\tau}}(g_e) \geq l_e^{\hat{\tau}}(g_e)$ for all e (part (i) of Lemma 3.9), this implies that

$$\begin{aligned} x_{j,R} \left(\sum_{e \in R} (\tilde{l}_e^{\hat{\tau}}(g_e) + \hat{\tau}_e) - \sum_{e \in Q} (\tilde{l}_e^{\hat{\tau}}(g_e) + \hat{\tau}_e) \right) &\leq x_{j,R} \left((1 + \gamma) l_R^{\hat{\tau}}(g) - l_Q^{\hat{\tau}}(g) \right) + 2\epsilon D \\ &\leq (1 + \gamma) x_{j,R} (l_R^{\hat{\tau}}(g) - l_Q^{\hat{\tau}}(g)) + \gamma D l_Q^{\hat{\tau}}(g) + 2\epsilon D \\ &\leq (1 + \gamma) \epsilon D + \gamma D \cdot 2mU + 2\epsilon D \leq 6\epsilon D. \end{aligned}$$

The penultimate inequality above follows since $l_e^*(g_e) + \hat{\tau}_e \leq 2U$ for every edge e . Let $g_e = \alpha_e a_e + (1 - \alpha_e) b_e$, where $\alpha_e \in [0, 1]$, and a_e, b_e are consecutive values in X . Note that the α_e s have encoding length $\text{poly}(\mathcal{I}, \log(\frac{1}{\delta}))$. Then the inequality

$$x_{j,R} \left(\sum_{e \in R} (\alpha_e l_{e,a_e} + (1 - \alpha_e) l_{e,b_e} + \hat{\tau}_e) - \sum_{e \in Q} (\alpha_e l_{e,a_e} + (1 - \alpha_e) l_{e,b_e} + \hat{\tau}_e) \right) \leq 6\epsilon D$$

yields a separating hyperplane of encoding length $\text{poly}(\mathcal{I}, \frac{1}{\sigma}, \log(\frac{1}{\delta}))$. ■

3.2 Extensions

Linear constraints on tolls given by a separation oracle. Here, we require that the tolls τ^* imposing the target flow f^* should lie in some polyhedron X , where X is given by means of a separation oracle. This is rich enough to model the following interesting scenarios.

- A subset F of edges cannot be tolled. This corresponds to the explicit constraint $\tau_e = 0$ for all $e \in F$.
- The total toll paid by any player under the flow f^* is at most a given budget B . This corresponds to the constraints $\tau(P) \leq B$ for every commodity i and path $P \in \mathcal{P}^i$ with $f_P^{*i} > 0$. One can separate over these exponentially-many constraints efficiently via a longest-path computation since f^* is acyclic.

The only change to our algorithm is that we first check if our current toll-vector $\hat{\tau}$ lies in X . If not then the separation oracle provided yields the separating hyperplane; otherwise, we proceed as before. The query complexity is now polynomial in the input size and the encoding length of X .

Nonatomic routing games with heterogeneous players. In a routing game $\Gamma = (G, l, \mathcal{K})$ with heterogeneous players, each commodity $i \in [k]$ has additionally a parameter $\gamma \in \mathbb{R}_+$. For flow f and tolls τ , a player of commodity i has delay $l_e(f_e) + \gamma_i \tau_e$ on edge e (for homogenous players, $\gamma_i = 1$ for each commodity i). In our query model, we assume we are given G , the commodity set $\mathcal{K} = \{(s_i, t_i, d_i, \gamma_i)\}$, and the feasible target multicommodity flow f^* .

General nonatomic congestion games. This is a generalization of network routing games, where the graph is replaced by an arbitrary set E of resources, and $\mathcal{P}^i \subseteq 2^E$ is the strategy-set associated with player-type i ; a more complete definition appears in Appendix B. Our ellipsoid-based algorithm uses essentially no information about the underlying graph. We only require that given a congestion-vector f , we can find the maximum-delay set $P \in \mathcal{P}^i$ for a given player-type i , and can find a valid decomposition of f of small support. Both of these are trivial since the \mathcal{P}^i sets are explicitly given in the input. Thus, our algorithms readily extend to general nonatomic congestion games and Theorems 3.2 and 3.6 (with mk replaced by $\sum_i |\mathcal{P}^i|$) continue to hold.

Atomic splittable routing games. Here, each commodity i represents a *single* player who controls d_i volume of flow and her strategy is to choose an s_i - t_i flow f^i of value d_i . The cost incurred by a player i under a feasible multicommodity flow (i.e., strategy profile) $f = (f^i)_{i \leq k}$ is $\sum_e f_e^i l_e(f_e)$.

Our results extend to atomic splittable routing games if we assume that for all valid choices of parameters of the latency functions and tolls (as encountered during the ellipsoid method), the underlying atomic splittable routing game has a unique Nash equilibrium. Here, by uniqueness we mean that if f and g are two Nash equilibria, then $f_e^i = g_e^i$ for all commodities i and edges e . This is not without loss of generality, but is known to hold, for example, if all latency functions are convex polynomials of degree at most 3, or if the graph is a generalized nearly-parallel graph and $x l_e(x)$ is strictly convex for all e (see [6]). When we say that tolls τ induce a flow $f^* = (f^{*i})_{i \leq k}$ here, we mean that the flow of every commodity i on every edge

e is f_e^{*i} in the resulting equilibrium. Our result shows that the task of computing tolls that induce specific *commodity-flows* can be reduced to the task of computing Nash equilibria (under the uniqueness assumption), even in the black-box setting. Although, to our knowledge, no algorithm is known for either of these tasks, even when latency functions are given, we believe that this reduction is of independent interest. The proof of Theorem 3.12 is very similar to that of Theorem 3.2: the only change is that to find the separating hyperplane, we now consider the marginal delay functions instead of the delay functions; see Appendix B.

Theorem 3.12. *In an atomic splittable routing game satisfying the aforementioned assumption, tolls that induce a target flow $f^* = (f^{*i})_{i \leq k}$ at equilibrium, if they exist, can be obtained with a polynomial number of queries to an oracle that returns the equilibrium flow under tolls.*

3.3 An algorithm for series-parallel networks with near-linear query complexity

We now give an algorithm for series-parallel networks with $\tilde{O}(m)$ query complexity. This is a significant improvement over the ellipsoid-based algorithm, and almost matches the linear lower bound proved in Theorem 5.1 for single-commodity routing games on parallel-link graphs with linear latency functions.

Theorem 3.13. *On two-terminal series-parallel graphs, one can compute in polynomial time tolls that induce a given target multicommodity flow f^* using $\tilde{O}(m)$ queries to an oracle that returns the equilibrium flow. Thus, we obtain $\tilde{O}(m)$ query complexity for multicommodity routing games with standard linear delay functions.*

We first recall some relevant details about series-parallel graphs. A *two-terminal directed series-parallel graph*, abbreviated series-parallel (sepa) graph, with terminals s and t is defined inductively as follows. A basic sepa graph is a directed edge (s, t) . Given two sepa graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, with terminals s_1, t_1 and s_2, t_2 respectively, one can create a new sepa graph $G = (V, E)$ as follows. A *series join* of G_1 and G_2 yields the graph obtained by identifying t_1 and s_2 , with terminals $s = s_1$ and $t = t_2$. A *parallel join* of G_1 and G_2 yields the graph obtained by identifying s_1 and s_2 , and t_1 and t_2 ; its terminals are $s = s_1 = s_2$ and $t = t_1 = t_2$.

For every series-parallel graph $G = (V, E)$, the recursive construction naturally yields a binary *decomposition tree*. The leaves of the tree are edges of G , and each internal node specifies a series- or a parallel-join. Each node of the tree also represents a subgraph of the G (obtained by performing the joins specified by the subtree rooted at that node), which is also clearly a sepa graph. In the sequel, we fix a decomposition tree corresponding to G . Whenever we say a subgraph of G , we mean a subgraph corresponding to a node of this decomposition tree. Given a subgraph H , we use s_H, t_H to denote its two terminals, and $\mathcal{P}(H)$ to denote the set of all s_H - t_H paths. We sometimes call s_H and t_H , the source and sink of H respectively. Let \mathcal{H} be the collection of subgraphs corresponding to the parallel-join nodes of the decomposition tree. For each $H \in \mathcal{H}$ obtained via the parallel join of H_1 and H_2 , we identify one of these as the “left” subgraph H_L and the other as the “right” subgraph H_R . Let \mathcal{P} denote the set of all s - t paths, where $s = s_G, t = t_G$.

Proof outline. Before we delve into the proof of Theorem 3.13, we give some intuition and give a roadmap of the proof. It is useful to first consider the simplest case of a graph with two parallel edges. Observe that any target flow can be obtained by varying the *difference* in tolls on these two edges. Further, the correct difference in tolls can be obtained by a binary search. Our key insight is that this intuition can be extended to series-parallel graphs via a suitable transformation of tolls. We show that tolls required to obtain a target flow can actually be described by the difference in tolls for each pair of parallel subgraphs, and then use binary search to obtain the correct differences that yield the target flow.

Formally, we show that any edge tolls in a sepa graph can in fact be transformed into certain canonical tolls that are defined in terms of subgraphs (Claim 3.15). Further, formalizing the intuition that what is

relevant is only the difference in tolls on parallel subgraphs, we make the novel connection that canonical tolls are in fact equivalent to labels on subgraphs $H \in \mathcal{H}$ (Lemma 3.16), where the label on subgraph $H \in \mathcal{H}$ stores the difference in the canonical tolls of subgraphs H_L and H_R whose parallel-join yields H .

Thus, our problem reduces to finding the correct labels on subgraphs $H \in \mathcal{H}$, which we aim to find via binary search. To do so, we establish certain structural properties of multicommodity flows in sepa graphs (Lemma 3.18). We leverage these to argue that if the canonical edge-tolls obtained from our current labels do not enforce the target flow, then we can find a subgraph $H \in \mathcal{H}$ and deduce whether its label should be increased or decreased. The query complexity is thus at most $|\mathcal{H}|$ times a logarithmic term depending on the accuracy required and the parameters of the routing game. A detailed description appears after Claim 3.19.

The presence of multiple commodities complicates things, since in the particular decomposition tree that we fix for G , all edges in a subgraph may be shortest-path edges for one commodity but not for another. This creates problems with the binary search since Claim 3.19 may not hold. We handle this by first arguing that there always exist tolls enforcing f^* such that every s - t path, and hence every s_i - t_i path is a shortest-path under edge costs $(l_e^* \tau^*(f^*))_e$ (Claim 3.14).

We believe that our structural insights into tolls and multicommodity flows on sepa graphs are of independent interest and likely to find other applications. In fact, our results on flows in sepa graphs also play an important role in our algorithm for inducing target flows via Stackelberg routing in Section 4.

Claim 3.14. *For $\Gamma^* = (G, l^*, (s_i, t_i, d_i)_{i \leq k})$ and target flow f^* there exist tolls $\tau^* \in \mathbb{R}_+^E$ such that:*

- (i) $\min_{P \in \mathcal{P}} \tau^*(P) = 0$;
- (ii) $l_P^*(f^*) + \tau^*(P) = l_Q^*(f^*) + \tau^*(Q)$ for every i and paths $P, Q \in \mathcal{P}^i$; and therefore
- (iii) $f(l^*, \tau^*) = f^*$.

Proof. We will show that for any edge costs $(c_e)_e$, there exist tolls τ so that every s - t path is a shortest path under edge costs $(c_e + \tau_e)_e$, and $\min_{P \in \mathcal{P}} \tau(P) = 0$. The claim follows simply by taking edge costs $(c_e = l_e^*(f^*))_e$ and setting $\tau^* = \tau$, since every s_i - t_i path clearly belongs to some s - t path.

The proof is by induction on the height of the decomposition tree for G . In the base case, if the decomposition tree has height 1, G consists of a single edge and setting $\tau_e = 0$ satisfies the claim. For the inductive step, suppose G is formed by the composition of H_1 and H_2 , and let c^1 and c^2 be the edge costs in subgraphs H_1 and H_2 respectively. Let τ^1 and τ^2 be the tolls that satisfy the claim for costs c^1 in subgraph H_1 , and costs c^2 in subgraph H_2 respectively.

If G consists of H_1 and H_2 composed in series, let $\tau_e = \tau_e^1$ if $e \in E(H_1)$ and $\tau_e = \tau_e^2$ otherwise. Then since any s - t path P consists of an s_1 - t_1 path and an s_2 - t_2 path, each of which is a shortest path in H_1 and H_2 respectively, every s - t path is a shortest path. Secondly, by the inductive hypothesis, there is a path P in H_1 with $\tau^1(P) = 0$, and a path Q in H_2 with $\tau^2(Q) = 0$. The concatenation of paths P and Q yields an s - t path R with $\tau(R) = 0$.

Suppose G consists of H_1 and H_2 composed in parallel. For any paths $P \in \mathcal{P}(H_1)$ and $Q \in \mathcal{P}(H_2)$, let $\delta = c(Q) + \tau^2(Q) - c(P) - \tau^1(P)$. We may assume that $\delta \geq 0$ (otherwise switch H_1 and H_2). Note that by the inductive hypothesis the value of δ is independent of the choice of P and Q . Define tolls τ for graph G as follows:

$$\tau_{vw} = \begin{cases} \tau_{vw}^1 + \delta, & \text{if } v = s \text{ and } (v, w) \in E(H_1). \\ \tau_{vw}^1, & \text{if } v \neq s \text{ and } (v, w) \in E(H_1). \\ \tau_{vw}^2, & \text{if } (v, w) \in E(H_2). \end{cases}$$

Then for any s - t path P , if $P \in \mathcal{P}(H_1)$ then $c(P) + \tau(P) = c(P) + \tau^1(P) + \delta$. If $Q \in \mathcal{P}(H_2)$ then $c(Q) + \tau(Q) = c(Q) + \tau^2(Q)$. By definition of δ and the induction hypothesis, every s - t path is

thus a shortest s - t path. Since the tolls on paths in H_2 remain the same, there is also an s - t path R with $\tau(R) = 0$. ■

Claim 3.15. For any tolls $\tau \in \mathbb{R}_+^E$ on the edges of G , there exist $\alpha \in \mathbb{R}_+^E$ such that:

(i) $\tau(P) = \alpha(P)$ for all $P \in \mathcal{P}$, and

(ii) for every subgraph H and every edge $e = (s_H, v) \in E(H)$, $\alpha_e \geq \min_{P \in \mathcal{P}(H)} \alpha(P)$.

Proof. The proof is again by induction on the height of the decomposition tree. If G is a single edge, then $\alpha = \tau$. If G is composed of subgraphs H_1 and H_2 , let τ^1 and τ^2 be the projection of τ onto the subgraphs. If H_1 and H_2 are in parallel, and tolls α^1 and α^2 satisfy the claim for tolls τ^1 and τ^2 in the subgraphs, it is easy to verify that tolls α defined by $\alpha_e = \alpha_e^1$ for $e \in E(H_1)$ and $\alpha_e = \alpha_e^2$ for $e \in E(H_2)$ satisfy the claim.

If H_1 and H_2 are in series, let α^1 and α^2 satisfy the claim for tolls τ^1 and τ^2 in the subgraphs. Define $\delta = \min_{P \in \mathcal{P}(H_2)} \alpha^2(P)$ and define the tolls

$$\alpha_{vw} = \begin{cases} \alpha_{vw}^1 + \delta, & \text{if } v = s_1 \text{ and } (v, w) \in E(H_1) \\ \alpha_{vw}^1, & \text{if } v \neq s_1 \text{ and } (v, w) \in E(H_1) \\ \alpha_{vw}^2 - \delta, & \text{if } v = s_2 \text{ and } (v, w) \in E(H_2) \\ \alpha_{vw}^2, & \text{if } v \neq s_2 \text{ and } (v, w) \in E(H_2). \end{cases}$$

Any s - t path P consists of segment P_1 between vertices $s = s_1$ and t_1 , and segment P_2 between $t_1 = s_2$ and $t = t_2$. Then

$$\alpha(P) = \alpha(P_1) + \alpha(P_2) = \alpha^1(P_1) + \delta + \alpha^2(P_2) - \delta = \tau^1(P_1) + \tau^2(P_2) = \tau(P).$$

Thus the first part of the claim is satisfied.

For the second part, consider any subgraph H . If $H = G$, then since every path $P \in \mathcal{P}(H)$ consists of segments $P_1 \in \mathcal{P}(H_1)$ and $P_2 \in \mathcal{P}(H_2)$, for every edge $e = (s, v) \in E$,

$$\begin{aligned} \alpha_{sv} &= \alpha_{sv}^1 + \delta \\ &\geq \min_{P \in \mathcal{P}(H_1)} \alpha^1(P) + \min_{P \in \mathcal{P}(H_2)} \alpha^2(P) && \text{(by the inductive hypothesis and definition of } \delta) \\ &= \min_{P \in \mathcal{P}(H_1)} \alpha^1(P) + \delta + \min_{P \in \mathcal{P}(H_2)} \alpha^2(P) - \delta \\ &= \min_{P \in \mathcal{P}(H)} \alpha(P). \end{aligned}$$

If $H \neq G$, then since every path $P \in \mathcal{P}(H)$ contains exactly one edge incident to s_H , the toll along every path changes by exactly the same quantity ($+\delta$, $-\delta$, or zero). ■

We call tolls $\alpha \in \mathbb{R}_+^E$ that satisfy property (ii) of Claim 3.15 *canonical tolls*. Thus, any edge tolls can be modified to obtain canonical edge tolls α . These in turn can be mapped to a *labeling* (L, Δ) , where $\Delta = (\Delta_H)_{H \in \mathcal{H}} \in \mathbb{R}_+^{\mathcal{H}}$, by setting $L = \min_{P \in \mathcal{P}} \alpha(P)$, and $\Delta_H = \min_{P \in \mathcal{P}(H_L)} \alpha(P) - \min_{P \in \mathcal{P}(H_R)} \alpha(P)$ for all $H \in \mathcal{H}$. Lemma 3.16 shows that this mapping is in fact invertible. Given the labeling (L, Δ) we can obtain canonical edge tolls α by the following procedure. Note that $|\mathcal{H}| \leq m$.

-
- M1. Initialize $\alpha_e = 0$ for all e .
 - M2. We traverse subgraphs in \mathcal{H} in a bottom-up manner, i.e., we consider all subgraphs in \mathcal{H} that are descendants of $H \in \mathcal{H}$ before considering H . When we consider a subgraph H , we set $\alpha_e = \alpha_e + \max\{0, \Delta_H\}$ for all $e = (s_H, v) \in E(H_L)$, and $\alpha_e = \alpha_e + \max\{0, -\Delta_H\}$ for all $e = (s_H, v) \in E(H_R)$.
 - M3. Finally, we set $\alpha_e = \alpha_e + L$ for all $e = (s, v) \in E$.
-

Lemma 3.16. *Let (L, Δ) be the labeling obtained from some canonical tolls $\alpha \in \mathbb{R}_+^E$, and β be the tolls obtained from (L, Δ) by the above procedure. Then $\alpha = \beta$.*

Proof. Let β' be the tolls obtained after step 1 of the above procedure, i.e., before adding L to the edges incident to s . We will show that for each edge e not incident to s , $\beta'_e = \alpha_e$, while for each edge e incident to s , $\beta'_e = \alpha_e - \min_{P \in \mathcal{P}} \alpha(P)$.

The proof is by induction on the size of G . If $G = \{e\}$, then since there are no parallel compositions, $\mathcal{H} = \emptyset$, and hence $\beta'_e = 0 = \alpha_e - \min_{P \in \mathcal{P}} \alpha(P)$. If G is the series-join of H_1 and H_2 , then for each edge not incident to s_{H_1} or s_{H_2} , $\beta'_e = \alpha_e$ by the inductive hypothesis. Further, note that the minimum toll $\alpha(P)$ over all s_{H_2} - t_{H_2} paths must be zero, since otherwise, on any edge $e = (s, v)$, α_e would be strictly less than the minimum toll over s - t paths. Hence by the inductive hypothesis $\beta'_e = \alpha_e$ for edges that leave s_{H_2} . For edges incident to s , since any s - t path consists of a path between s and $t_1 = s_2$ and between s_2 and t , and by the inductive hypothesis,

$$\beta'_e = \alpha_e - \min_{P \in \mathcal{P}(H_1)} \alpha(P) = \alpha_e - \min_{P \in \mathcal{P}} \alpha(P)$$

where the second equality follows because, as earlier observed, the minimum toll $\alpha(P)$ over all s_{H_2} - t_{H_2} paths must be zero. Thus the inductive hypothesis holds in this case.

If G is the parallel-join of H_1 and H_2 , then for each edge not incident to s , $\beta'_e = \alpha_e$ by the inductive hypothesis. Further, assume without loss of generality that $\min_{P \in \mathcal{P}(H_1)} \alpha(P) = \min_{P \in \mathcal{P}} \alpha(P)$. Then by the inductive hypothesis, for each edge $e = (s, v) \in E(H_1)$,

$$\beta'_e = \alpha_e - \min_{P \in \mathcal{P}(H_1)} \alpha(P) = \alpha_e - \min_{P \in \mathcal{P}} \alpha(P)$$

as stated in the claim. Let $\delta = \min_{P \in \mathcal{P}(H_2)} \alpha(P) - \min_{P \in \mathcal{P}(H_1)} \alpha(P) \geq 0$. By the procedure for computing β' , if $H_1 = H_L$ and $H_2 = H_R$, then $\Delta_H = -\delta$, otherwise $\Delta_H = \delta$. In both cases, when considering G , we only modify the tolls on edges of $E(H_2)$ incident to s_{H_2} by adding δ to these. So for each edge $e = (s, v) \in E(H_2)$, we have

$$\begin{aligned} \beta'_e &= \alpha_e - \min_{P \in \mathcal{P}(H_2)} \alpha(P) + \delta \\ &= \alpha_e - \min_{P \in \mathcal{P}(H_2)} \alpha(P) + \min_{P \in \mathcal{P}(H_2)} \alpha(P) - \min_{P \in \mathcal{P}(H_1)} \alpha(P) \\ &= \alpha_e - \min_{P \in \mathcal{P}} \alpha(P) \end{aligned}$$

which completes the induction step, and hence, the proof. ■

Definition 3.17. Given multicommodity flows f and \tilde{f} , we call a pair H_1, H_2 of subgraphs, (f, \tilde{f}) -discriminating if:

- (i) the parallel-join of H_1 and H_2 is a subgraph in \mathcal{H} ; and
- (ii) $f_e > \tilde{f}_e$ for all $e \in E(H_1)$, and $f_e \leq \tilde{f}_e$ for all $e \in E(H_2)$.

Lemma 3.18. *Let f and \tilde{f} be two feasible multicommodity flows for $(G, (s_i, t_i, d_i)_{i \leq k})$. If $f \neq \tilde{f}$, then there exists an (f, \tilde{f}) -discriminating pair of subgraphs.*

Proof Sketch. We use induction on the series-parallel structure to first show a slightly weaker statement: there exist subgraphs H_1 and H_2 whose parallel join is in \mathcal{H} such that: (a) $f_e \geq \tilde{f}_e$ for all $e \in E(H_1)$, $f_e \leq \tilde{f}_e$ for all $e \in E(H_2)$, and (b) $|f_{H_1}|$, which we define to be the total flow routed under f in H_1 for commodities not internal to H_1 , is greater than $|\tilde{f}_{H_1}|$, and $|f_{H_2}| < |\tilde{f}_{H_2}|$. Now if $f_e > \tilde{f}_e$ for all $e \in E(H_1)$

then we are done. Otherwise, we show that if we consider the minimal subgraph K of H_1 (under the same decomposition tree used for G) that contains both $f_e > \tilde{f}_e$ and $f_e = \tilde{f}_e$ edges, then K must be a parallel-join of subgraphs that form an (f, \tilde{f}) -discriminating pair. ■

We defer a full proof of Lemma 3.18 until Appendix C.

Claim 3.19. *Let $\hat{f} = f(l^*, \tau)$. If there is a subgraph H such that $\hat{f}_e > f_e^*$ for all $e \in E(H)$ then there is some commodity i such that every s_H - t_H path is part of a shortest s_i - t_i path under edge costs $(l_e^{*\tau}(\hat{f}_e))_e$.*

Proof. The proof is by induction on the size of H . If H is an edge e , there is some commodity i such that $\hat{f}_e^i > 0$, so the statement holds. If H is the parallel join of H_1, H_2 , then it follows from the induction hypothesis that every s_H - t_H path must be of equal length (since there are commodities corresponding to both H_1 and H_2); hence, there is a commodity corresponding to H and the statement follows. Suppose H is the series composition of H_1, H_2 . Let \mathcal{K} be the set of commodities i such that $\sum_{e=(s_H, v) \in E(H)} \hat{f}_e^i > 0$. For every $i \in \mathcal{K}$ such that $t_i \in V(H) \setminus \{t_H\}$, the set of edges $(s_H, v) \in E(H)$ forms an s_i - t_i cut, and so the flow across the cut must be the same in \hat{f}^i and f^{*i} . However, $\sum_{e=(s_H, v) \in E(H)} \hat{f}_e > \sum_{e=(s_H, v) \in E(H)} f_e^*$, so there is some commodity $j \in \mathcal{K}$ such that $s_j, t_j \notin V(H) \setminus \{s_H, t_H\}$. For commodity j , some s_H - t_H path is part of a shortest s_j - t_j path under edge costs $(l_e^{*\tau}(\hat{f}_e))_e$. Applying the induction hypothesis to H_1, H_2 yields that all s_H - t_H paths are of the same length. Thus, every s_H - t_H path is part of a shortest s_j - t_j paths under edge costs $(l_e^{*\tau}(\hat{f}_e))_e$. ■

We now describe the algorithm for Theorem 3.13. Let τ^* be tolls given by part (b) of Claim 3.14 and $(0, \Delta^*)$ be the labeling obtained from τ^* . We may assume that $\tau_e^* \in [0, U']$ and is a multiple of $\frac{1}{U'}$ for all e , where $U' = m \text{poly}(U, \sum_i d_i)$. E.g., with standard linear latencies, since every $f_e^*, a_e^*, b_e^* \in [0, U]$ and is a multiple of $\frac{1}{U}$, we can take $U' = \max\{U^2, mK \sum_i d_i\}$.

-
- T1. Initialize, $L_H = -mU', U_H = mU', \Delta_H = 0$ for all $H \in \mathcal{H}$. Let $L = 0$. Let $M = m \log(8mU'^2)$.
- T2. For $r = 1, \dots, M$, we do the following. Map (L, Δ) to canonical tolls α as described in steps M1–M3. Query the oracle to obtain $\hat{f} = f(l^*, \alpha)$. If $\hat{f} = f^*$, then exit the loop. Otherwise, find an (\hat{f}, f^*) -discriminating pair of subgraphs H_1, H_2 (which exists by Lemma 3.18).
Let H be the parallel join of H_1, H_2 . If $H_1 = H_L$, update $L_H \leftarrow \Delta_H$, else update $U_H \leftarrow \Delta_H$. If $|U_H - L_H| < \frac{1}{U'}$, set Δ_H to be the multiple of $\frac{1}{U'}$ in $[L_H, U_H]$; else update $\Delta_H = (L_H + U_H)/2$.
- T3. Return tolls α .
-

Proof of Theorem 3.13. Let α^* be the canonical tolls obtained from τ^* via Claim 3.15, and let (L^*, Δ^*) be the corresponding labeling. We have $L^* = 0$ due to Claims 3.14 and 3.15. The proof of Claim 3.15 shows that, under the assumptions on τ^* , we have α_e^* is a multiple of $\frac{1}{U'}$, and is in $[0, mU']$ for all e . Hence, $\Delta_H^* \in [-mU', mU']$ and is a multiple of $\frac{1}{U'}$, for all $H \in \mathcal{H}$.

We say that the intervals $[L_H, U_H]$ assigned to $H \in \mathcal{H}$ are valid if $\Delta_H^* \in [L_H, U_H]$ for all $H \in \mathcal{H}$. We argue below that our algorithm maintains valid intervals. Give this, in each iteration we halve the length of some interval, and this may happen at most $\log(8mU'^2)$ times for the interval of some $H \in \mathcal{H}$ until we find Δ_H^* , since Δ_H^* is a multiple of $\frac{1}{U'}$. Since there are at most m subgraphs in \mathcal{H} , after M iterations (without reaching f^*), we obtain Δ^* .

We now prove that the algorithm maintains valid intervals. Given tolls τ and a subgraph H , define $\tau_H := \min_{P \in \mathcal{P}(H)} \tau(P)$. So $\Delta_H^* = \alpha_{H_L}^* - \alpha_{H_R}^*$. The intervals are clearly valid at the start of the algorithm. Suppose the intervals are valid at the start of an iteration in step T2. We may assume that $\hat{f} \neq f^*$. By Claim 3.19, there is some commodity i such that every s_{H_1} - t_{H_1} path is part of a shortest s_i - t_i path under edge costs $(l_e^*(\hat{f}_e) + \alpha_e)_e$. Let $P = \operatorname{argmin}_{P' \in \mathcal{P}(H_1)} \alpha^*(P')$ and $Q = \operatorname{argmin}_{Q' \in \mathcal{P}(H_2)} \alpha(Q')$. Since P is a segment of a shortest-path for commodity i , We have

$$l_P^*(f^*) + \alpha_{H_1} < l_P^*(\hat{f}) + \alpha_{H_1} \leq l_P^*(\hat{f}) + \alpha(P) \leq l_Q^*(\hat{f}) + \alpha(Q) = l_Q^*(\hat{f}) + \alpha_{H_2} \leq l_Q^*(f^*) + \alpha_{H_2}.$$

Here, the first and last inequalities follow since H_1, H_2 is (\hat{f}, f^*) -discriminating. The second inequality follows from the definition of α_{H_1} ; the third, since P is part of a shortest s_i - t_i path; and the fourth equality, from the definition of Q . We know that every s - t path is a shortest s - t path under edge costs $(l_e^*(f_e^*) + \alpha_e^*)_e$. So we have

$$l_P^*(f^*) + \alpha_{H_1}^* = l_P^*(f^*) + \alpha^*(P) = l_Q^*(f^*) + \alpha^*(Q) \geq l_Q^*(f^*) + \alpha_{H_2}^*.$$

Combining this with the earlier inequality gives $\alpha_{H_1} - \alpha_{H_2} < \alpha_{H_1}^* - \alpha_{H_2}^*$. So if $H_1 = H_L$, then $\Delta_H < \Delta_H^*$; otherwise, $\Delta_H > \Delta_H^*$. Thus, our update for H preserves the validity of the intervals. \blacksquare

Remark 3.20. Our analysis shows that the above algorithm works whenever we have a “sign oracle” that given input tolls τ and a flow f^* , returns the sign of $f(l^*, \tau)_e - f_e^*$ for all edges e . This is clearly weaker than having an exact-equilibrium oracle.

Tolls in sepa graphs with an approximate oracle We focus on single-commodity routing games on sepa graphs. As before, we use \mathcal{H} to denote the collection of parallel subgraphs of G under a given sepa decomposition tree for G . For any tolls τ we use $\tilde{f}(\tau)$ to denote the flow returned by the ϵ -oracle, and use $f(\tau)$ to denote the actual equilibrium flow. Note that $\tilde{f}(\tau)$ may not be a continuous function of τ . In fact, if queried twice with the same tolls, the oracle may return different flows. However, on any edge, $|\tilde{f}_e(\tau) - f_e(\tau)| \leq \epsilon$, and hence for any subgraph H , $|\tilde{f}_H(\tau) - f_H(\tau)| \leq m\epsilon$. For any tolls τ , we define $N(\tau) := \max_P \tau(P) - \min_P \tau(P) + ml_{\max}(d)$.

Here, we describe the algorithm and the main theorem. The proof of the theorem is in Appendix F.

We now describe the algorithm for computing tolls τ' so that $f(\tau')$ is close to f^* on every edge.

-
- T1. We traverse all subgraphs $H \in \mathcal{H}$ of G in a bottom-up manner, i.e., we consider all subgraphs of H before considering H . For each subgraph in this order, let τ be the current tolls. Since $H \in \mathcal{H}$, let H be the parallel composition of H_1 and H_2 . Let $LB = -N(\tau)$, $UB = N(\tau)$.
- T2. Let $\gamma = 0$ initially, and $\gamma_1 = (LB+UB)/2$. By binary search in $[-2N(\tau), 2N(\tau)]$, find γ so that $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$ where τ' is as given by Lemma F.3. If $|\tilde{f}_{H_1}(\tau')| - |f_{H_1}^*| \geq 6m\epsilon$, set $LB = \gamma_1$ and repeat this step. If $|\tilde{f}_{H_1}(\tau')| - |f_{H_1}^*| \leq -6m\epsilon$, set $UB = \gamma_1$ and repeat this step. Otherwise, let $\tau = \tau'$ and continue with the next subgraph.
-

Theorem 3.21. *The algorithm terminates with at most $m^3 \log^2(20ml_{\max}(d)/\kappa(\epsilon/m))$ queries and with tolls τ' so that $|f_e(\tau') - f_e^*| \leq 14m^2\epsilon$ on every edge.*

3.4 Nearly quadratic query complexity for single-commodity, linear-delay routing games

For the case of linear delays, we show that our bound on the query-complexity from Section 3.1 can be significantly improved upon. We give an algorithm with quadratic query complexity even in general graphs (ignoring logarithmic terms in the query complexity). Interestingly, the quadratic complexity is used to obtain tolls that give an equilibrium flow with the same support — i.e., on the same edges — as the target flow. Following this, actually obtaining the correct tolls requires only a linear number of queries.

Theorem 3.22. *For a single-commodity routing game Γ with standard linear delay functions, tolls that enforce f^* can be obtained in $\tilde{O}(m^2)$ queries.*

Throughout, we assume without loss of generality that $f^* > 0$; otherwise, we impose infinite tolls on any edge where $f_e^* = 0$, effectively removing these edges from the graph.¹ We assume the delay function on any edge e is $l_e(x) = a_e x + b_e$. Define $l_{\max}(x) := \max_{e \in E} a_e x + b_e$, and $\kappa(x) = x^2/Kd$. Define the

¹The use of infinite tolls is a notational convenience; the same effect can be obtained with tolls $m^2 2^m l_{\max}(d)$.

support of a flow f to be the set of edges with strictly positive flow. We will use negative tolls in our proof; however, by Claim 3.23 which we prove in Appendix D, this is again just a notational convenience. Similar arguments were used in [15] to show boundedness of tolls, but the results are not directly applicable. Note that f^* is acyclic.

Claim 3.23. *For a single-commodity routing game and tolls τ , there exist tolls $\tau' \geq 0$ so that $f(\tau) = f(\tau')$ and $\tau'_{e'} \leq \tau_{e'} + \sum_{e:\tau_e < 0} |\tau_e|$ for all e' . If the graph is acyclic, τ' can be obtained without knowledge of the delay functions.*

Proof outline. We show that if the support of the equilibrium flow remains fixed, the equilibrium flow is a linear function of the tolls. Thus if we can obtain tolls τ so that the support of $f(\tau)$ is the same as f^* , we can solve a linear system of equations to obtain tolls that enforce f^* . Accordingly, our algorithm consists of the following two steps.

Step 1: Enforcing the correct support. We first obtain tolls τ so that $f_e(\tau) > 0 \Leftrightarrow f_e^* > 0$. By suitably large tolls on edges e for which $f_e^* = 0$, we already have tolls that satisfy one direction of the implication. The other direction is roughly by binary search, described in Lemma 3.26: we pick an edge r that does not yet have flow, and impose increasingly negative tolls on this edge until it has positive flow at the equilibrium. The difficulty here is in maintaining monotonicity of the support of the equilibrium flow. Increasing the flow on edge r decreases flow on the other edges. We use a number of results regarding the sensitivity of equilibrium flow for this step. In fact, this step has quadratic query complexity, while the second step that actually obtains tolls that enforce f^* has linear query complexity.

Step 2: Obtaining the target flow f^* . We now use Lemma 3.28 which establishes the linearity of equilibrium flow as a function of tolls, if the support of the equilibrium flow does not change. Obtaining the coefficients of this linear map requires us to query the oracle with a small toll on each edge. The query complexity of this step is thus linear.

We start with some results about the continuity, monotonicity, and sensitivity of equilibrium flow as a function of tolls. Theorem 3.24(ii) was earlier proved in [13]. Let $\mathbb{1}_e \in \mathbb{R}^E$ be the vector with value 1 in coordinate e , and 0 everywhere else.

Theorem 3.24. *Let Γ be a single-commodity routing game with standard linear delay functions. Then,*

- (i) $\|f(0) - f(\mathbb{1}_e \kappa(\epsilon))\|_\infty \leq \epsilon$,
- (ii) $f(\tau)$ is continuous,
- (iii) for edge r and $\delta > 0$, $f_r(\mathbb{1}_r \delta) \leq f_r(0)$, and
- (iv) for edge r and $\delta > 0$, $|f_r(-\mathbb{1}_r \delta) - f_r(0)| \geq \|f(-\mathbb{1}_r \delta) - f(0)\|_\infty$.

The proof of (i) and (ii) are straightforward from the following immediate Corollary of Lemma 3.5. We prove (iii) and (iv) in Appendix D.

Corollary 3.25. *For a multicommodity routing game Γ , let \hat{f} be the equilibrium flow. If g is a valid multicommodity flow that satisfies for all i , $P \in \mathcal{P}^i$, $g_P > 0 \Rightarrow l_P(g) \leq D^i(l, g) + \epsilon$, then $\|g - \hat{f}\|_\infty \leq \sqrt{K \epsilon \sum_i d_i}$.*

We now show a lemma that is used to prove the first step of our proof. Lemma 3.26 shows that if edge r has no flow or very little flow at equilibrium, then with a small number of queries we can obtain tolls so that the flow on edge r increases, and the flow on the other edges does not change significantly.

Lemma 3.26. *Let Γ be a single-commodity nonatomic routing game, and let $\delta > 0$, $\delta \leq d$. For tolls τ , let $S := \{e : f_e(\tau) \geq \delta\}$, and edge $r \notin S$. Then with $\log(-N/\kappa(\delta/3))$ queries, we can determine tolls τ' so that $f_e(\tau') \geq \delta/3$ for all $e \in S \cap \{r\}$, where $N := \min_{P \in \mathcal{P}} \tau_P - \min_{P \in \mathcal{P}: r \in P} \tau_P - ml_{\max}(d) < 0$.*

Proof. To obtain tolls τ' , we will only vary the tolls on edge r . We thus parametrize tolls τ' by α , where $\tau' = \tau + \mathbf{1}_r \alpha$.

If $f_r(\tau) \geq \delta/3$, we are done. Otherwise, we claim that if $\alpha = N$, then $f_r(\tau') = d$. To see this, let Q be the path that minimizes $\sum_{e \in P} a_e d + b_e + \tau_e$ over all paths $P \in \mathcal{P}$ with $r \in P$, and let f be the flow that sends the entire demand along this path. Then the delay along this path with tolls τ' is

$$\begin{aligned} \sum_{e \in Q} (a_e d + b_e + \tau'_e) &= \sum_{e \in Q} (a_e d + b_e + \tau_e) + N \\ &= \min_{P: r \in P} \sum_{e \in P} (a_e d + b_e + \tau_e) + \min_{P \in \mathcal{P}} \tau(P) - \min_{P \in \mathcal{P}: r \in P} \tau(P) - ml_{\max}(d) \\ &\leq \min_{P \in \mathcal{P}} \tau(P), \end{aligned}$$

while for any path P with $r \notin P$, the delay along path P is at least this quantity. Hence f is actually an equilibrium flow, and if $\alpha \leq N$, then $f_r(\tau') = d$.

Define $a, b \in [N, 0]$ as follows.

$$\begin{aligned} a &:= \max\{\alpha \in [N, 0] : f_r(\tau') = \delta/3\} \\ b &:= \min\{\alpha \in [N, 0] : f_r(\tau') = 2\delta/3\} \end{aligned}$$

By the continuity of equilibrium flow with respect to tolls (Theorem 3.24, (ii)), a, b exist. By the monotonicity of equilibrium flow, for any $\alpha \in [b, a]$, $f_r(\tau') \in [\delta/3, 2\delta/3]$. Then by Theorem 3.24, (iv), for any edge $e \in S$ and $\alpha \in [b, a]$, $f_e(\tau') \in [\delta/3, \delta]$. Thus our problem reduces to finding an $\alpha \in [b, a]$, which we can find by binary search. We will show that $a - b \geq \kappa(\delta/3)$, which gives us the bound on the number of queries required. To see this, let τ^a and τ^b be the tolls obtained by setting $\alpha = a$ and $\alpha = b$ respectively. Then if $a - b \leq \kappa(\delta/3)$, then by Theorem 3.24 (i), $\|f(\tau^a) - f(\tau^b)\| \leq \delta/3$. \blacksquare

F1. Initialize $\tau_e \leftarrow 0$ for all e , $i \leftarrow 1$, and $S \leftarrow \{e : f_e(\tau) \leq d/3^i\}$.

F2. While $S \neq E$

F3. Pick an edge $r \notin S$

F4. By Lemma 3.26, find $\alpha \in [N, 0]$ so that if $\tau' = \tau + \mathbf{1}_r \alpha$, then $f_e(\tau') \geq d/3^{i+1}$ for all $e \in S \cup \{r\}$.

F5. $\tau \leftarrow \tau'$, $i \leftarrow i + 1$, $S \leftarrow \{e : f_e(\tau) \leq d/3^i\}$

Lemma 3.27. *The stated algorithm terminates with tolls τ so that $f_e(\tau) \geq d/3^m$ on every edge, and requires $O(m^2 \log(3ml_{\max}(d)))$ queries.*

Proof of Lemma 3.27. Let $N(i)$ be the value of N in the i th iteration of the while loop. Then by Lemma 3.26, the i th iteration requires $\log(-N(i)/\kappa(d/3^{i+1}))$ queries to complete, and adds at least one edge to the set S . Thus, there are at most m iterations of the while loop. We will show that $|N(i)| \leq m2^{i-1}l_{\max}(d)$, thus proving the bound on the number of queries. Note that since all tolls are negative, $|\min_P \tau_P - \min_{P: r \in P} \tau_P| \leq |\min_P \tau_P|$.

The proof is by induction. In the first iteration since $\tau = 0$ initially, $N(1) \leq ml_{\max}(d)$. In the i th iteration, there are at most $i - 1$ other edges with tolls on them, and along any path the sum of the absolute values of these tolls is at most $\sum_{j \leq i-1} 2^{j-1} ml_{\max}(d) = (2^{i-1} - 1)ml_{\max}(d)$, and hence $|N(i)| \leq 2^{i-1} ml_{\max}(d)$. ■

This completes the first step of our proof. We now proceed with the second step. Lemma 3.28 shows that the equilibrium flow is a linear function of the tolls, as long as the set of edges with strictly positive flow remains constant. While a similar result on the linearity of the equilibrium flow was shown in [11], Lemma 3.28 shows how to obtain the coefficients of the linear map.

Lemma 3.28. *For any routing game Γ and tolls $\tau^{(1)}$, let $f(\tau^{(1)}) > 0$. Then there exist coefficients $(\beta_{e,e'})_{e,e' \in E}$ so that for any tolls τ ,*

$$(i) \quad f(\tau + \tau^{(1)}) > 0 \Rightarrow f(\tau + \tau^{(1)}) = f(\tau^{(1)}) + \beta\tau, \text{ and}$$

$$(ii) \quad f(\tau^{(1)}) + \beta\tau > 0 \Rightarrow f(\tau + \tau^{(1)}) = f(\tau^{(1)}) + \beta\tau.$$

Proof. We first show how to obtain the coefficients $(\beta_{e,e'})_{e,e'}$. Define $f_{\min} = \min_e f(0) > 0$. For each edge e' , let $\alpha^{e'} := \mathbb{1}_{e'} \kappa(f_{\min}/2)$. By Corollary 3.25, $f(\tau^{(1)} + \alpha^{e'}) > f_{\min}/2$ for each edge e' . Then for each edge $e \in E$, define $\beta_{e,e'} = \left(f_e(\tau^{(1)} + \alpha^{e'}) - f_e(\tau^{(1)}) \right) / \kappa(f_{\min}/2)$.

Given tolls τ , let $g := f(\tau^{(1)}) + \sum_{e'} \beta_{e,e'} \tau_{e'}$. In general, g may be negative on some edges. However, we show that g is an s - t pseudoflow of value d : it satisfies all the conditions for being a flow except nonnegativity. Further, we show that g is a minimizer of (1) if we allow each f to be a pseudoflow, rather than a flow.

To see the first claim, note that for a fixed edge e' since $\beta_{e,e'}$ is the difference of two (scaled) flows of the same value, it is a circulation. Then g is the sum of a flow and a set of circulations, and is hence a pseudoflow.

To show that g equalizes the delay on every s - t path with tolls τ , for any s - t path p ,

$$\begin{aligned} \sum_{e \in p} l_e(g) - l_e(f(\tau^{(1)})) &= \sum_{e \in p} a_e (g_e - f_e(\tau^{(1)})) = \sum_{e \in p} a_e \sum_{e'} \beta_{e,e'} \tau_{e'} \\ &= \sum_{e'} \frac{\tau_{e'}}{\kappa(f_{\min}/2)} \sum_{e \in p} a_e \left(f_e(\tau^{(1)} + \alpha^{e'}) - f_e(\tau^{(1)}) \right) \end{aligned} \quad (6)$$

Since $f(\tau^{(1)} + \alpha^{e'})$ and $f(\tau^{(1)})$ are equilibrium flows with tolls $\tau^{(1)} + \alpha^{e'}$ and $\tau^{(1)}$ respectively, and both are strictly positive on every edge, it follows from (6) that

$$\sum_{e \in p} l_e(g) - l_e(f(\tau^{(1)})) = \sum_{e'} \frac{\tau_{e'}}{\kappa(f_{\min}/2)} \left(D(f(\tau^{(1)} + \alpha^{e'})) - D(f(\tau^{(1)})) - \sum_{e \in p} \alpha_e^{e'} \right)$$

and since $\alpha_e^{e'} = 0$ for $e \neq e'$,

$$\sum_{e \in p} l_e(g) - l_e(f(\tau^{(1)})) = \sum_{e'} \frac{\tau_{e'}}{\kappa(f_{\min}/2)} \left(D(f(\tau^{(1)} + \alpha^{e'})) - D(f(\tau^{(1)})) \right) - \sum_{e \in p} \tau_e.$$

Thus for any path p ,

$$\sum_{e \in p} l_e(g) + \tau_e = \sum_{e \in p} l_e(f(\tau^{(1)})) + \sum_{e'} \frac{\tau_{e'}}{\kappa(f_{\min}/2)} \left(D(f(\tau^{(1)} + \alpha^{e'})) - D(f(\tau^{(1)})) \right).$$

Further, for any path p , $\sum_{e \in p} l_e(f(\tau^{(1)})) = D(f(\tau^{(1)} - \sum_{e \in p} \tau_e^{(1)}))$. Hence for any path p , $\sum_{e \in p} l_e(g) + \tau_e + \tau_e^{(1)}$ is equal. It follows immediately that if the second condition in the lemma is true, i.e., if $g > 0$, then g must be an equilibrium flow with tolls $\tau^{(1)} + \tau$, and since the equilibrium is unique, $f(\tau + \tau^{(1)}) = g$. This completes the proof of the second statement.

For the first statement, for $0 \leq \lambda \leq 1$ define $h(\lambda) = f(\tau^{(1)} + \tau) + \lambda(g - f(\tau^{(1)} + \tau))$. Since on any path p as shown earlier $\sum_{e \in p} l_e(g) + \tau_e + \tau_e^{(1)}$ is equal, and $f(\tau^{(1)} + \tau) > 0$ by assumption, this is also true for $h(\lambda)$. Further since $f(\tau^{(1)} + \tau) > 0$, there exists $\lambda > 0$ so that $h(\lambda) > 0$. Then $h(\lambda)$ must also be an equilibrium flow with tolls τ . By the uniqueness of equilibria, this is only possible if $f(\tau + \tau^{(1)}) = g$. ■

L1. Use the earlier algorithm to get tolls $\tau^{(1)}$ so that $f(\tau^{(1)}) \geq d/2^m$.

L2. Obtain the coefficients $(\beta_{e,e'})_{e,e'}$ as in Lemma 3.28

L3. Solve the linear equations $\beta\tau^{(2)} = f^* - f(\tau^{(1)})$ for tolls $\tau^{(2)}$. Then $f(\tau^{(2)} + \tau^{(1)}) = f^*$.

Proof of Theorem 3.22. We will show that the algorithm is correct, and requires $O(m^2 \log(3ml_{\max}(d)))$ queries. The correctness of the first step follows from Lemma 3.27. To use Lemma 3.28, since $f_{\min} \geq d/2^m$, to obtain the coefficients $(\beta_{e,e'})_{e,e'}$, we require an additional m queries, each of which applies an additional toll (relative to $\tau^{(1)}$) of $\kappa(d/2^{m+1})$ on individual edges.

Let τ^* be tolls such that $f(\tau^*) = f^*$. By the first part of Lemma 3.28, then $f^* = f(\tau^{(1)}) + \beta(\tau^* - \tau^{(1)})$. Now $\tau^{(2)}$ is a solution to the system of linear equalities $\beta\tau^{(2)} = f^* - f(\tau^{(1)})$; since $\tau^* - \tau^{(1)}$ satisfies this, we know a solution exists. Further, by the second part of the Lemma since $f(\tau^{(1)}) + \beta\tau^{(2)} = f^* > 0$, in fact $f(\tau^{(1)} + \tau^{(2)}) = f^*$. Hence $\tau^{(1)} + \tau^{(2)}$ are the tolls required to obtain the target flow. ■

4 Inducing target flows via Stackelberg routing on series-parallel graphs

Recall that here we have a single-commodity routing game $\Gamma^* = (G, l^*, (s, t, d))$. We are given a parameter $\alpha \in [0, 1]$ and a target flow f^* , and we seek an s - t flow g of value of at most αd such that $g + f(l^*, g) = f^*$, if one exists. We abbreviate $f(l^*, g)$ to $f(g)$. We consider the setting where G is a directed sepa graph with terminals s and t , and devise an efficient algorithm that computes a Stackelberg routing inducing f^* using at most m queries to an oracle that returns the equilibrium flow. The flow g we compute is in fact of minimum value among all Stackelberg flows that induce f^* . (So either g is the desired Stackelberg flow, or none exists if $|g| > \alpha d$.) Our algorithm works for arbitrary increasing delay functions provided, as in Section 3.3, we have an oracle that returns the correct sign of $((f(g) + g)_e - f^*)_e$ given a Stackelberg routing g . In particular, the algorithm works for increasing linear latencies.

As before, we fix a decomposition tree for G , and a subgraph refers to a subgraph corresponding to a node of this tree. For a flow f and subgraph H , let f_H denote $(f_e)_{e \in E(H)}$. We again leverage the concept of a good pair of subgraphs, which becomes much simpler to state in the single-commodity setting.

Definition 4.1 (specialization of Definition C.1). Given s - t flows f, \tilde{f} , we call a pair of subgraphs H_1, H_2 (f, \tilde{f}) -good if:

- (i) the parallel-join of H_1, H_2 is a subgraph;
- (ii) $f_e \geq \tilde{f}_e$ for all $e \in E(H_1)$ and $f_e \leq \tilde{f}_e$ for all $e \in E(H_2)$; and

(iii) $|f_{H_1}| > |\tilde{f}_{H_1}|$ and $|f_{H_2}| < |\tilde{f}_{H_2}|$.

Lemma 4.2. *Let g be any Stackelberg routing. If $f(g) + g \neq f^*$, there exists an $(f(g) + g, f^*)$ -good pair of subgraphs.*

Lemma 4.2 follows from a more general result proved in Lemma C.2 for multicommodity flows. The proof in the single-commodity setting becomes much simpler, and follows immediately from Claim 4.3 since $|f(g) + g| = |f^*|$.

Claim 4.3. *For any two s - t flows f, \tilde{f} in a sepa graph G , either there is an (f, \tilde{f}) -good pair of subgraphs, or one of the following holds:*

- (i) *If $|f| = |\tilde{f}|$ then $f = \tilde{f}$.*
- (ii) *If $|f| > |\tilde{f}|$ then $f \geq \tilde{f}$.*
- (iii) *If $|f| < |\tilde{f}|$ then $f \leq \tilde{f}$.*

Proof. The proof is by induction on the size of the graph. For a single edge, there is no good pair of subgraphs, but one of the three cases must hold. For the induction step, let G be the join of subgraphs G_1 and G_2 . Let $f_1 = f_{G_1}$, $\tilde{f}_1 = \tilde{f}_{G_1}$, and $f_2 = f_{G_2}$, $\tilde{f}_2 = \tilde{f}_{G_2}$. Clearly, f_1, \tilde{f}_1 are s_{G_1} - t_{G_1} flows, and f_2, \tilde{f}_2 are s_{G_2} - t_{G_2} flows. If G_1 contains an (f_1, \tilde{f}_1) -good pair, or G_2 contains an (f_2, \tilde{f}_2) -good pair, then the same pair is an (f, \tilde{f}) -good pair, and we are done. So assume otherwise.

Suppose G_1 and G_2 are in series. Then, $|f_1| = |f| = |f_2|$, and $|\tilde{f}_1| = |\tilde{f}| = |\tilde{f}_2|$. So whichever case applies to f and \tilde{f} , the same applies to f_1, \tilde{f}_1 , and f_2, \tilde{f}_2 . By the induction hypothesis, we have the desired relationship between f_1, \tilde{f}_1 and f_2, \tilde{f}_2 , and hence between f and \tilde{f} . So the statement holds for G .

Suppose G_1 and G_2 are in parallel. If $|f_1| > |\tilde{f}_1|$ and $|f_2| < |\tilde{f}_2|$, then by the induction hypothesis, $f_1 \geq \tilde{f}_1$, $f_2 \leq \tilde{f}_2$, so G_1, G_2 is an (f, \tilde{f}) -good pair. Similarly, if $|f_1| < |\tilde{f}_1|$ and $|f_2| > |\tilde{f}_2|$, then G_2, G_1 is an (f, \tilde{f}) -good pair. So assume neither case holds.

Now if $|f| = |\tilde{f}|$, then (after eliminating the above cases) $|f_1| = |\tilde{f}_1|$, $|f_2| = |\tilde{f}_2|$. Hence, by the induction hypothesis, we have $f_1 = \tilde{f}_1$, $f_2 = \tilde{f}_2$, and so $f = \tilde{f}$.

If $|f| > |\tilde{f}|$, then it must be that $|f_1| \geq |\tilde{f}_1|$ and $|f_2| \geq |\tilde{f}_2|$. Therefore, $f_1 \geq \tilde{f}_1$, $f_2 \geq \tilde{f}_2$, and so $f \geq \tilde{f}$.

Finally, if $|f| < |\tilde{f}|$, then it must be that $|f_1| \leq |\tilde{f}_1|$, $|f_2| \leq |\tilde{f}_2|$. Hence, $f_1 \leq \tilde{f}_1$, $f_2 \leq \tilde{f}_2$, and so $f \leq \tilde{f}$. This completes the induction step, and hence, the proof. \blacksquare

Our algorithm is now quite simple to describe. We keep track of the set \bar{S} , initialized to \emptyset , of edges not on any shortest s - t path under the edge costs $(l_e^*(f_e^*))_e$. By Lemma 2.2, \bar{S} must be saturated by any Stackelberg routing that induces f^* . We repeatedly do the following.

-
- S1. Find the flow g of minimum value that saturates every edge in \bar{S} and satisfies $g_e \leq f_e^*$ for all e .
 - S2. Query the oracle with g as the Stackelberg flow. If $f^* = f(g) + g$, exit and return g . Otherwise, find an $(f(g) + g, f^*)$ -good pair of subgraphs H_1, H_2 . Add every edge in H_2 to \bar{S} (and repeat the process).
-

Theorem 4.4. *The above algorithm computes a Stackelberg flow g of minimum value that induces f^* in at most m queries.*

Proof. In every iteration, $|\bar{S}|$ increases by at least 1: since $|f_{H_2}(g) + g_{H_2}| < |f_{H_2}^*|$ and g saturates every edge in \bar{S} , we know that at least one edge in H_2 is not in the current set \bar{S} . When $\bar{S} = E$, we have $g = f^*$. So the algorithm terminates in at most m iterations with some flow g that induces f^* . To complete the proof, we only need to show that any edge added to \bar{S} is indeed a non-shortest-path edge. Let $h = f(g) + g$. Let $s' = s_{H_1} = s_{H_2}$, $t' = t_{H_1} = t_{H_2}$. Since $|h_{H_1}| > |f_{H_1}^*|$, there is some s' - t' path P in H_1 such that $(h - f^*)_P > 0$. So P belongs to a shortest s - t path under edge costs $(l_e^*(h_e))_e$. So for every s' - t' path Q in H_2 , we have $l_P^*(f^*) < l_P^*(h) \leq l_Q^*(h) \leq l_Q^*(f^*)$. So every edge of H_2 is a non-shortest-path edge under edge costs $(l_e^*(f_e^*))_e$. \blacksquare

Stackelberg routing with an approximate oracle Define an ϵ -oracle for Stackelberg flows (or simply an ϵ -oracle) for routing game (G, l, s, t, d, \cdot) as an oracle that takes as input a Stackelberg flow g and returns an ϵ -equilibrium of the routing game $(G, l(g; \cdot), s, t, d - |g|)$ having encoding length $\text{poly}(\mathcal{I}, \log(1/\epsilon))$. In this section, we show that we can use such an ϵ -oracle to obtain a Stackelberg flow g that approximately enforces f^* , and has value at most that of g^* , i.e., the Stackelberg flow that enforces f^* . Define $\epsilon' = \sqrt{Kd\epsilon}$. As before, we fix a particular decomposition of the series-parallel graph G . Each subgraph then corresponds to a node in the decomposition tree.

Here, we give the main definition, the algorithm, and the theorem statement. The proof of the theorem is in Appendix E.

Definition 4.5. Given s - t flows f, \tilde{f} , we call a pair of subgraphs H_1, H_2 (f, \tilde{f}) - ϵ' -good if, with $m' := |E(H_1)| + |E(H_2)|$,

1. The parallel-join of H_1, H_2 is a subgraph,
2. There exists $r \in E(H_1)$ with $f_r - \tilde{f}_r \geq (2mK)^{2m'} \epsilon'$,
3. For every subgraph H' of H_1 with $r \in E(H')$, $|f_{H'}| - |\tilde{f}_{H'}| \geq 2(m+1)\epsilon'$,
4. For every edge $e \in E(H_2)$, $f_e - \tilde{f}_e \leq (2mK)^{2|E(H_2)|} \epsilon'$, and
5. $|f_{H_2}| - |\tilde{f}_{H_2}| < 0$.

Our algorithm is similar to the algorithm given for the case of the exact oracle, except that instead of a good pair of subgraphs, we find an ϵ' -good pair of subgraphs. We keep track of the set \bar{S} , initialized to \emptyset , of edges not on any shortest s - t path under the edge costs $(l_e(f_e^*))_e$. By Lemma 2.2, \bar{S} must be saturated by any Stackelberg routing that induces f^* . We repeatedly do the following.

-
- S1. Find the flow g of minimum value that saturates every edge in \bar{S} and satisfies $g_e \leq f_e^*$ for all e .
 - S2. Query the ϵ -oracle with g as the Stackelberg flow and let \tilde{f} be the flow returned by the oracle. If $\|\tilde{f} + g - f^*\|_\infty \leq (2mK)^{2m} \sqrt{Kd\epsilon}$, exit and return g . Otherwise, find an $(\tilde{f} + g, f^*)$ - ϵ' -good pair of subgraphs H_1, H_2 . Add every edge in H_2 to \bar{S} (and repeat the process).
-

Theorem 4.6. Let g^* be a Stackelberg flow that induces f^* . Then with m ϵ -oracle queries, the above algorithm finds a Stackelberg flow g so that $\|f(g) + g - f^*\|_\infty \leq (2mK)^{2m} \epsilon' + \epsilon'$ and $|g| \leq |g^*|$.

5 Query- and computational- complexity lower bounds

Our algorithms in the previous section show that tolls for enforcing a target flow can be obtained with a polynomial number of queries in general, and for routing games with linear latencies, with near-linear and near-quadratic queries for series-parallel and general graphs respectively. For Stackelberg routing, we similarly give a linear-query algorithm for series-parallel graphs. We now give a lower bound for the query-complexity with tolls, showing that at least a linear number of queries are required in the worst case to enforce a target flow. The bound is in fact for parallel links with linear delay functions, and hence our algorithm for series-parallel graphs is nearly optimal. We then consider the related problem of determining an equivalent set of delay functions, and give strong evidence that this is a hard problem. We show that even the computational question — i.e., determining if two given sets of delay functions are equivalent — is *NP*-hard, both for tolls and Stackelberg routing. For Stackelberg routing, we additionally give exponential lower bounds for the query-complexity of determining an equivalent set of delay functions. Finally, using the techniques developed in this section, we prove that given a routing game, the optimization problem of determining the Stackelberg routing to minimize cost of the remaining players, called the Stackelberg Equilibrium Delay Minimization problem, is *APX*-hard.

5.1 A linear lower bound for query complexity with tolls

We show a lower bound of $\Omega(m)$ on the number of queries required to obtain tolls that give the target flow.

Theorem 5.1. *Any deterministic algorithm that computes tolls required to enforce a target flow requires $\Omega(m)$ queries, even for a single commodity instance on parallel links with linear delay functions.*

Our example for the lower bound consists of a single commodity on m parallel links, with the demand $d = m$ and the target flow $f_e^* = 1$ on each edge. In fact, our lower bound is actually for the problem of obtaining tolls τ with the right support: $f_e(\tau) > 0$ iff $f_e^* > 0$.

For our lower bound example, our delay functions are defined by a permutation $\pi^* : [m] \rightarrow [m]$. The delay function on the parallel edge e_i is given by $(x/m) + 2(\pi^*(i) - 1)$. Thus, we use the notation $f(\pi, \tau)$ for the equilibrium flow, where the permutation π identifies the delay functions. We show that any algorithm that computes the correct tolls to enforce f^* must obtain the correct permutation π^* , and we design an oracle that after k queries has only revealed information about $\pi^{*-1}(1), \dots, \pi^{*-1}(k)$. Thus, in order to compute the correct tolls, any algorithm requires $m - 1$ queries.

Our oracle works as follows. Initially, let $A^0 = \emptyset$ be the set of assigned edges in the partial permutation π^* . For the j th query $\tau^j = (\tau_e^j)_{e \in E}$, our oracle returns the equilibrium flow described below.

Oracle: Pick an arbitrary edge e that with minimum toll τ_e^j , so that e is not in A^{j-1} . Let $\pi^*(e) = j$ and $A^j = A^{j-1} \cup \{e\}$. Let $\pi_{(j)}$ be a complete permutation that extends the partial permutation π^* , and return $f(\pi_{(j)}, \tau^j)$ as the equilibrium flow in response to tolls τ^j .

Claim 5.2. *For $j \in [m]$, let σ be any permutation that satisfies $\sigma(e) = \pi_{(j)}(e)$ for all edges $e \in A^j$. Then for any edge $e \notin A^j$, $f_e(\sigma, \tau^j) = 0$*

Proof. By description of the oracle and since $e \notin A^j$, e is not the unique edge with minimum toll in τ^j . That is, there exists an edge $h \in A^j$ with $\tau_h^j \leq \tau_e^j$. Since $h \in A^j$, $\sigma(h) \leq j < \sigma(e)$. Further since $\tau_h^j \leq \tau_e^j$, by description of the delay functions, $l_e(0) > l_h(d)$, and hence edge e cannot have flow at equilibrium. ■

We now show that the equilibrium flows returned by our oracle are consistent.

Lemma 5.3. *There exists a permutation σ so that for every $j \in [m]$, $f(\pi_{(j)}, \tau^j) = f(\sigma, \tau^j)$.*

Proof. Fix $j \leq m$, and let σ be a complete permutation that extends π^* . The image of every edge $e \in A^j$ is the same in π^* and $\pi_{(j)}$. Thus the delay functions on both edges is the same. By Claim 5.2, the equilibrium is zero for any edge not in A^j . Since edges in A^j have the same delay function, the equilibrium flow must be the same for permutations σ and $\pi_{(j)}$. ■

Proof of Theorem 5.1. From Lemma 5.3, for any sequence of $m - 1$ toll queries, the oracle returns a consistent sequence of equilibrium responses. Further, from Claim 5.2, since $|A^j| \leq m - 1$ for $j \leq m - 1$, there is an edge e with no flow in the equilibrium returned by the oracle. Hence, since f^* has positive flow on every edge, any deterministic algorithm requires at least m queries to compute tolls that obtain f^* . ■

5.2 Lower bounds for determining equivalence

Given the ability to query a routing game and obtain the equilibrium flow, a natural question is if we can in fact obtain the delay functions on the edges. It is obvious that the exact delay functions cannot be obtained, even for a single edge. However, is it possible to obtain delay functions that are equivalent, in the sense that any tolls or Stackelberg routing would yield the same equilibrium flow as in the routing game? In fact, we show that the problem of determining equivalence of delay functions is *NP*-hard, even if we are given two sets of delay functions and our task is to determine if they are equivalent. For Stackelberg routing we additionally show exponential lower bounds for the query complexity.

5.2.1 Equivalence with tolls

For tolls, we show that even with delay functions that are *constant* on every edge, determining equivalence is NP-hard. Since with constant delay functions we may have multiple equilibria, we instead consider the set of edges on any shortest path from s to t . For delay function $l = (l_e)_e$ and tolls τ on the edges we define

$$\text{sp}(l, \tau) := \{e : e \text{ lies on a shortest path from } s \text{ to } t \text{ with edge lengths } l(f(l, \tau)) + \tau\}.$$

Definition 5.4. Given a graph G with demand d between s and t , two sets of delay functions on the edges l^1 and l^2 are equivalent with regard to tolls if for every vector of tolls $(\tau_e)_e$, $\text{sp}(l^1, \tau) = \text{sp}(l^2, \tau)$.

Theorem 5.5. Given two constant cost functions $l^1 = (c_e^1)_e$ and $l^2 = (c_e^2)_e$, it is NP-hard to determine if there exist tolls $\tau \geq 0$ so that $\text{sp}(l^1, \tau) \neq \text{sp}(l^2, \tau)$.

The reduction for NP-hardness is from the problem of 2-Directed Disjoint Paths, which is known to be NP-complete [17]:

Definition 5.6 (2-Directed Disjoint Paths (2DDP)). Given a directed graph G and vertices s_1, s_2, t_1 and t_2 , do there exist s_i - t_i paths p_i such that p_1 and p_2 are vertex-disjoint?

Proof of Theorem 5.5. Given an instance of 2DDP on graph G , we modify the instance as follows. We add two vertices s and t . We add 6 additional edges. We first add edges (s, s_1) , (s, s_2) , (t_1, t) , and (t_2, t) . We also add edges (s, t) and (t_1, s_2) . This last edge is called the *diagonal* edge. This gives us graph $G' = (V, E')$.

In l^1 , every edge in G' has cost $l_e^1 = 1$. In l^2 , every edge except the diagonal edge has cost 1, while the diagonal edge has length 2.

We first show that if G has 2 directed disjoint paths, then there exist the required tolls. Let p_1, p_2 be the vertex-disjoint paths between s_1, t_1 and s_2, t_2 respectively, and let m_i be the number of edges in p_i . In graph G' , let

$$F := \{p_1 \cup p_2 \cup \{(s, s_1), (t_1, s_2), (t_2, t), (s, t)\}\}$$

and consider the tolls given by

$$\tau_e = \begin{cases} |F| & e \notin E' \\ m_1 + m_2 + 3/2, & e = (s, t) \\ 0 & \text{otherwise} \end{cases}$$

Then with cost function l^1 and tolls τ , the path along edge (s, t) is the shortest path of total cost $m_1 + m_2 + 5/2 = |F| - 3/2$, while with cost function l^2 and tolls τ the edges of $F \setminus \{(s, t)\}$ form the shortest path of total cost $|F| - 1$. Thus, $\text{sp}(l^1 + \tau) \neq \text{sp}(l^2 + \tau)$.

Now we will show that if there exist tolls τ so that $\text{sp}(l^1 + \tau) \neq \text{sp}(l^2 + \tau)$, there exist vertex-disjoint paths between s_1, t_1 and s_2, t_2 . Given such tolls, let $S_i := \text{sp}(l^1 + \tau)$. Since each edge has a cost of at least 1 and all tolls are nonnegative, S_1 and S_2 must be acyclic. Further, since $l_e^1 = l_e^2$ for every edge other than the diagonal edge, this edge must be in exactly one of the sets, say S_1 . Then S_1 must contain a simple s - t path p containing the diagonal edge. Since the diagonal edge is (t_1, s_2) , and by construction the only edges incident on s are (s, s_1) and (s, s_2) , p must consist of the (s, s_1) edge, followed by an s_1 - t_1 path, the diagonal edge, and s_2 - t_2 path, and the edge (t_2, t) . Since the path is simple, the s_i - t_i paths must be disjoint, completing the proof. ■

5.2.2 Equivalence with Stackelberg routing

We start by defining equivalence for Stackelberg routing.

Definition 5.7. Given a graph G with demand d between s and t and a Stackelberg demand fraction α , two sets of delay functions on the edges l^1 and l^2 are ϵ -equivalent with regard to Stackelberg routing if for every Stackelberg routing g with $|g| \leq \alpha d$, $\|f(l^1, g) - f(l^2, g)\|_\infty \leq \epsilon$.

We prove strong lower bounds for both the query complexity and the computational complexity. In fact, for the query complexity, the lower bound instance is a graph of constant size. The size of the input is determined by the demand d , and we show that although the size of the input is $O(\log d)$ any deterministic algorithm that determines ϵ -equivalence for a fixed ϵ must make $\Omega(\sqrt{d})$ queries. For computational hardness, we show that even if we are explicitly given affine delay functions l^1 and l^2 , determining $1/2$ -equivalence is *NP*-hard. Our proof for computational hardness builds upon a reduction given by Roughgarden [34].

Query complexity. We are now given a graph G with demand d between s and t , a Stackelberg demand fraction α , and a set of delay functions l^1 on the edges of G . In addition, we are given query access to a second set of delay functions l^2 . As before, each query consists of a Stackelberg routing g , and the response is the equilibrium flow $f(l^2, g)$. We show the following result.




Figure 1: Braess graph instance for proving hardness of equivalence determination.

Theorem 5.8. Any deterministic algorithm that determines ϵ -equivalence for $\epsilon \leq 1/16$ requires an exponential number of queries.

Our proof of the theorem is based on a particular property exhibited by the Braess graph shown in Figure 1: there exist demands $d_1 \leq d_2$ that depend on the parameters a and b so that for any demand $d < d_1$ and $d > d_2$ the set of shortest-path edges is the same, and differs from the set of shortest-path edges for any demand $d_1 \leq d \leq d_2$. This is formalized by the following claim.

Claim 5.9. For the routing game depicted in Figure 1, and any $d_1, d_2 \in \mathbb{R}_+$ with $d_2 > d_1 \geq 1$, there exist parameters a and b so that the equilibrium flow f on the Braess edge is strictly positive iff $d_1 < d < d_2$, where d is the demand being routed. Further, if $d_2 - d_1 \geq \sqrt{2(d_1 + d_2)}$, then $f_{uv} \geq 1/12$ for demand $d = (d_1 + d_2)/2$.

Proof. We choose $a = 1 + (d_1 + d_2)/2$ and $b = (d_1 d_2)/4$. Then for any d , consider the flow that routes $d/2$ on the s - u - t path and $d/2$ on the s - v - t path. It is easy to verify that this is the equilibrium flow if and only if $d \leq d_1$ or $d \geq d_2$. Given the symmetric delay functions, it is then apparent that for $d \in (d_1, d_2)$ the (u, v) edge must have strictly positive flow.

For the second part of the proof, let $\sigma = d_1 + d_2$, $\delta = d_2 - d_1$. Thus $a = 1 + \sigma/2$, $b = (\sigma^2 - \delta^2)/16$, and $d = \sigma/2 = a - 1$. Let $f_{su} = x$. Then by the symmetry of the delay functions $f_{sv} = f_{ut} = d - x$ and $f_{uv} = 2x - d$. Since we know for this demand $f_{uv} > 0$, and edge (s, v) has zero delay if $f_{sv} = 0$, $l_{su}(f) + l_{uv}(f) - l_{sv}(f) = 0$. Hence

$$0 = x^2 + x + b - a(d - x)$$

and solving for $f_{uv} = 2x - d$, and substituting the values of a , b and d yields

$$\begin{aligned} 2x - d &= \sqrt{(a+1)^2 + 4(ad-b)} - (a+1+d) = x^2 + (a+1)x + (b-ad) \\ &= \sqrt{(2+\sigma)^2 + \delta^2/4} - (2+\sigma). \end{aligned}$$

Using the fact that $\sqrt{1+x} \geq 1 + x/3$ for $|x| \leq 1$ by the Taylor expansion, we get

$$2x - d \geq \frac{\delta^2}{12(2+\sigma)}$$

Since $\delta \geq \sqrt{2\sigma} \geq \sqrt{2+\sigma}$ by assumption, this completes the proof. \blacksquare

Proof of Theorem 5.8. We demonstrate that on the Braess graph in Figure 1 with an additional (s, t) edge, demand $d > 8$, and where a, b have value $O(d^2)$, any algorithm requires $\Omega(\sqrt{d})$ queries to determine if two sets of delay functions l^1, l^2 are equivalent. Since the size of the input is $O(\log d)$, this would prove the lemma.

For delay functions l^1 that are explicitly given, $a = 1, b = \infty$ and $l_{st}^1 = \infty$. Delay functions l^2 also have $l_{st}^2 = \infty$ but different values for a and b , which are determined after seeing the queries. Let $g^i, i \leq k$ be the set of queries. We will show that if $k \leq \sqrt{d}$ then there exist a, b so that $f(l^1, g^i) = f(l^2, g^i)$ for all $i \leq k$, but there exists g so that $f_{uv}(l^2, g) \geq 1/12$. Since $l_{uv}^1 = \infty$, it must be that $f_{uv}(l^1, g) = 0$, and hence the two delay functions are distinct. Thus any algorithm that makes less than \sqrt{d} must fail to distinguish between these delay functions.

For any query g^i , our oracle returns the equilibrium flow $f(l^1, g^i)$. Now given g^i for $i \leq k \leq \sqrt{d}$, let $\alpha_1, \alpha_2 \in [1, d]$ be such that $\alpha_2 - \alpha_1 \geq \sqrt{d}$, and for all $i, d - g_{st}^i \notin [\alpha_1, \alpha_2]$. Since edge (s, t) has infinite delay, any flow on this edge must be Stackelberg flow. Hence we require α_1 and α_2 so that the total flow on the Braess graph is always outside the interval $[\alpha_1, \alpha_2]$, and $\alpha_2 - \alpha_1 \geq \sqrt{d}$. Since $k \leq \sqrt{d}$, such an interval must exist. We then select a and b as in Claim 5.9 to complete our definition of delay function l^2 .

It remains to show that for all $g^i, f(l^2, g^i) = f(l^1, g^i)$ for correctness of the oracle. Fix i , and let $d' = d - g_{st}^i$. Note that $d' \notin [\alpha_1, \alpha_2]$. Let g_1, g_2 and g_3 be the Stackelberg flow on paths $s-u-t, s-u-v-t$, and $s-v-t$ respectively. By our choice of l^2 , if $g_1 = g_2 = g_3 = 0$, then the equilibrium flow would split demand d' equally between the $s-u-t$ and $s-v-t$ paths, and hence

$$l_{sv}^2(d'/2) \leq b + l_{su}^2(d'/2). \quad (7)$$

We consider the following cases.

Case 1: Either g_1 or g_3 is strictly greater than $(d' - g_2)/2$. Suppose $g_1 > (d' - g_2)/2$. We claim that at equilibrium, the non-Stackelberg demand is entirely routed on the $s-v-t$ path, i.e., $f(g) = h$ where $h_{sv}(g) = h_{vt}(g) = d' - (g_1 + g_2 + g_3)$. To see this, note that $h_{sv} + g_{sv} < d'/2$, hence comparing with (7), delay on the $s-v-t$ path is less than the delay on the $s-u-v-t$. Further, $h_{sv} + g_{sv} < g_{ut}$, and $h_{vt} + g_{vt} < g_{su}$. By the symmetry of delay functions, the $s-v-t$ path is therefore the shortest path, and hence $h = f(g)$.

Case 2: Both g_1 and g_3 are at most $(d' - g_2)/2$. We claim that at equilibrium, $f_{su}(g) = f_{ut}(g) = d'/2 - (g_1 + g_2)$ and $f_{sv}(g) = f_{vt}(g) = d'/2 - (g_3 + g_2)$. Thus at equilibrium the remaining flow $d' - g_2$ is divided equally between the $s-u-t$ and $s-v-t$ paths, and again the edge (u, v) has no flow at equilibrium. To verify the claim, note that

$$f_{sv}(g) + g_{sv} \leq f_{su}(g) + g_{su} = d'/2 \text{ and } f_{ut}(g) + g_{ut} \leq f_{vt}(g) + g_{vt} = d'/2$$

and hence, comparing with (7),

$$l_{sv}^2(f(g) + g) \leq l_{sv}^2(f(g) + g) + b \text{ and } l_{ut}^2(f(g) + g) \leq l_{ut}^2(f(g) + g) + b.$$

It is further easy to see that, since the total flow on edges (s, u) , (v, t) is equal, and the total flow on edges (s, v) , (u, t) is equal,

$$l_{sv}^2(f(g) + g) + l_{vt}^2(f(g) + g) = l_{su}^2(f(g) + g) + l_{ut}^2(f(g) + g).$$

Hence paths $s-u-t$ and $s-v-t$ are shortest paths with the described flow, and since $f_P(g) > 0$ only on these paths, it is an equilibrium.

As noted earlier, if the Stackelberg flow is rational, then so is $f(g)$. In fact as shown the equilibrium flow in all cases is very simple and can be computed directly from g . ■

We note that in our example, the equilibrium flow returned by the oracle is particularly simple and in fact does not depend on the delay functions. E.g., in the simpler case in the proof sketch, the oracle always returns $f_e(g) = (d - g_{st})/2$ for all $e \neq (s, t), (u, v)$.

Computational complexity. We now show that even if delay functions l^1 and l^2 are given explicitly, determining if they are ϵ -equivalent is computationally hard for $\epsilon \leq 1/2$. This is true even if all delay functions are affine. Our proof uses properties of the Braess graph together with ideas from a reduction from 2-Directed Disjoint Paths shown by Roughgarden [34].




Figure 2: Braess graph instance for proving hardness of equivalence determination with respect to Stackelberg routing.




Figure 3: 2DDP instance with additional edges for proving hardness of equivalence determination with respect to Stackelberg routing.

Theorem 5.10. *The problem of determining the ϵ -equivalence of delay functions for $\epsilon \leq 1/2$ is NP-hard.*

We use the following claim about Stackelberg routing in the Braess graph.

Claim 5.11. *In any Stackelberg routing instance on the graph with delay functions l as in Figure 2 and Stackelberg routing g , if $f_{uv}(g) > 0$, then $d < 2$ and $D(l, g) \leq 2$. Further, if $g_e = 0$ for every edge in the Braess graph and the demand $d \in [1/2, 3/2]$, then $f_{uv}(g) \geq 1/2$.*

Proof. For the first part of the claim, let $h = f(l, g) + g$. If $(u, v) \in \mathcal{S}(l, g)$, then the path $p = (s, u, v, t)$ must be a shortest path for flow h . Then $l_{su}(h) + l_{uv}(h) \leq l_{sv}(h)$, and hence $l_{su}(h) \leq 1$. Thus $h_{su} \leq 1$. Similarly, $h_{vt} \leq 1$. The first part follows. For the second part, if $d \leq 1$ it is easy to see that the equilibrium flow routes the entire demand on the s - u - v - t path. If $d \in [1, 2]$ then consider the flow $h_{su} = h_{vt} = 1$, $h_{uv} = 2 - d$ and $h_{sv} = h_{ut} = d - 1$. It can be verified that h is the equilibrium flow. ■

Proof of Theorem 5.10. We show a reduction from 2DDP. Given an instance of the 2DDP problem, after the addition of a source s and a sink t and additional edges described next (and shown in Figure 3), we add this graph in parallel with a standard Braess graph (Figure 2). The delay functions l^1, l^2 will differ only on edge (u, v) in the Braess graph. We use H_1 to refer to the Braess graph and H_2 to refer to the graph in the 2DDP instance with vertices s and t and the additional edges, and H to refer to their parallel composition. For a flow f , $|f_{H_i}|$ is the value of the flow in subgraph H_i .

The specifics of the construction are as follows. Let $m = |E|$ be the number of edges in the given instance of 2DDP. All of these edges have delay function x/m^2 . We add a source s , vertex s' and a sink t . We add an edge (s, s') with constant delay function $1/8$, and edges (s', s_i) and (t_i, t) for $i \in \{1, 2\}$. Edges (s, s_1) and (t_2, t) have delay function x , while edges (s, s_2) and (t_1, t) have delay function 1 . Further, there is an (s, t) edge with delay function ∞ , and for every edge $e = (x, y)$ in the original instance, the new instance additionally contains edges (s, x) and (y, t) with delay function ∞ . This constitutes the graph H_2 . Graph H_1 consists of the Braess graph instance in Figure 2, and graph H is obtained by a parallel composition of H_1 and H_2 . The Stackelberg instance has demand $m^4 + 3$, and $\alpha = m^4/(3 + m^4)$. The delay functions l^1, l^2 are as described, except $l_{uv}^2 = \infty$ on the Braess edge.

Since l^1, l^2 differ only on the delay function on edge (u, v) , it is easy to see that for any Stackelberg routing g , if $f_{uv}(l^1, g) = 0$ then $f_{uv}(l^1, g) = f_{uv}(l^2, g)$. Further, since $l_{uv}^2 = \infty$, if $f_{uv}(l^1, g) = f_{uv}(l^2, g)$ then in fact $f_{uv}(l^1, g) = 0$. Hence the delay functions are equivalent iff $f_{uv}(l^1, g) = 0$ for every Stackelberg routing g . For the proof of the theorem, we will show that if the instance of 2DDP is a positive instance, then there exists a Stackelberg routing g so that $f_{uv}(l^1, g) \geq 1/2$, otherwise for any Stackelberg routing, $f_{uv}(l^1, g) = 0$.

In the remainder of the proof we focus on delay functions l^1 . Suppose that the instance is a positive instance. Then the Stackelberg routing g sends m^3 units of flow on every edge $e = (x, y)$ in the original instance that is not on the vertex-disjoint paths p_i , using the additional edges $(s, v), (v, w)$. Any remaining Stackelberg flow is routed on the (s, t) edge. Thus every edge that is not on the vertex-disjoint paths now has delay at least m , while $g_e = 0$ for every edge on the vertex-disjoint paths. Further, $g_e = 0$ for every edge $e \in H_1$.

We claim that for the equilibrium flow, $1/2 \leq |f_{H_1}| \leq 3/2$. To see this, if $|f_{H_1}| < 1/2$, then the delay at equilibrium in H_1 is at most 1. However $|f_{H_2}| > 5/2$, hence at least one of the two s - t parallel paths has delay at equilibrium greater than 1. If $|f_{H_1}| > 3/2$ then the delay at equilibrium in H_1 is 2. However, $|f_{H_2}| \leq 3/2$, hence at least one of the two s - t parallel paths has delay at equilibrium less than $1 + 1/8 + 3/4 \times (1 + 1/m) < 2$. Thus at equilibrium, $1/2 \leq |f_{H_1}| \leq 3/2$, and by Claim 5.11, $f_{uv}(g) \geq 1/2$.

Now suppose that for some Stackelberg routing g , $f_{uv}(g) > 0$. Then by Claim 5.11, $|f_{H_1}| \leq 2$ and the delay at equilibrium is at most 2. However, then $|f_{H_2}| \geq 1$ and the delay at equilibrium is at most 2. Since there is an (s, s') edge with constant delay $1/8$, following the proof of Theorem 5.13, this is only possible if the instance of DDP is a positive instance. ■

In fact, using very similar ideas, we can show that the problem of minimizing $D(f(g))$ over all Stackelberg strategies is $(4/3 - \epsilon)$ -inapproximable, even with linear delays. Roughgarden has shown that finding the Stackelberg routing that minimizes the average delay of the *total flow* $g + f(g)$ is NP-hard, even in parallel links with affine delays [33]. Despite considerable interest in Stackelberg routing, nothing stronger than NP-hardness is known for this problem. Our result thus shows that a closely related problem is APX-hard.

Definition 5.12 (Stackelberg Equilibrium Delay Minimization (SEDM)). Given a Stackelberg routing instance $(G, l, (d, s, t), \alpha)$, find the Stackelberg routing g that minimizes the average delay for the equilibrium flow $f(g)$.

Theorem 5.13. *The SEDM problem is $(4/3 - \epsilon)$ -inapproximable, for any fixed $\epsilon > 0$.*

Proof. Given an instance of the 2DDP problem, we modify it to obtain a Stackelberg routing instance as follows. Let $m = |E|$ be the number of edges in the original instance. All of these edges have delay function x/m^2 . We add a source s and a sink t , and edges (s, s_i) and (t_i, t) . Edges (s, s_1) and (t_2, t) have delay function x , while edges (s, s_2) and (t_1, t) have delay function 1. Further, there is an (s, t) edge with delay function ∞ , and for every edge $e = (v, w)$ in the original instance, the new instance additionally contains edges (s, v) and (w, t) with delay function ∞ . The Stackelberg instance has demand $m^4 + 1$, and $\alpha = m^4/(1 + m^4)$.

We claim that if the instance of 2DDP is a positive instance, then there exists a Stackelberg routing g with $D(l, f(g)) \leq 3/2 + 1/m$, otherwise for any Stackelberg routing, $D(f(g)) \geq 2$. Suppose that the instance is a positive instance. Then the Stackelberg routing g sends m^3 units of flow on every edge $e = (v, w)$ in the original instance that is not on the vertex-disjoint paths p_i , using the additional edges (s, v) , (v, w) . Thus every edge that is not on the vertex-disjoint paths now has delay at least m , while $g_e = 0$ for every edge on the vertex-disjoint paths. Any remaining Stackelberg flow is routed on the (s, t) edge. It is now easy to verify that the equilibrium flow $f(g)$ splits one unit of demand approximately equally between the two paths p_1 and p_2 , and has a delay at equilibrium of at most $3/2 + 1/m$.

Now suppose the given instance does not contain two vertex-disjoint paths between s_1, t_1 and s_2, t_2 . Following the argument in [7], for a contradiction let g be a Stackelberg routing for which $D(f(g)) < 2$. Let F be the set of edges with positive flow at equilibrium. Then F must contain all four edges (s, s_1) , (s, s_2) , (t_1, t) , (t_2, t) ; the absence of any of these edges would give a delay of at least 2. Further, F cannot contain an s - s_2 - t_1 - t path since again this would give a delay of at least 2. Hence F must contain an s - s_1 - t_1 - t path and an s - s_2 - t_2 - t path. These paths cannot be vertex disjoint; let v be the common vertex. Then the delay on any s - v path must be at least 1, and the delay on any v - t path must be at least 1. Hence the total delay in any instance that does not contain two vertex-disjoint paths is at least 2, which gives us a contradiction. The hardness of determining the existence of these paths thus shows that minimizing the delay of the equilibrium flow with Stackelberg routing is $(4/3 - \epsilon)$ inapproximable, for any $\epsilon > 0$. ■

Acknowledgment. We thank Éva Tardos and Jim Geelen for useful discussions.

References

- [1] Kareem Amin, Rachel Cummings, Lili Dworkin, Michael Kearns, and Aaron Roth. Online learning and profit maximization from revealed preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 770–776, 2015.
- [2] Yakov Babichenko, Siddharth Barman, and Ron Peretz. Simple approximate equilibria in large games. In *EC*, pages 753–770, 2014.

- [3] Maria-Florina Balcan, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V. Vazirani. Learning economic parameters from revealed preferences. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, pages 338–353, 2014.
- [4] Martin Beckman, CB McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.
- [5] Xiaohui Bei, Wei Chen, Jugal Garg, Martin Hoefer, and Xiaoming Sun. Learning market parameters using aggregate demand queries. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 411–417, 2016.
- [6] Umang Bhaskar, Lisa Fleischer, Darrell Hoy, and Chien-Chung Huang. Equilibria of atomic flow games are not unique. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 748–757. Society for Industrial and Applied Mathematics, 2009.
- [7] Umang Bhaskar, Katrina Ligett, and Leonard J Schulman. The network improvement problem for equilibrium routing. In *17th IPCO*, pages 138–149, 2014. arXiv preprint arXiv:1307.3794.
- [8] Umang Bhaskar, Katrina Ligett, Leonard J. Schulman, and Chaitanya Swamy. Achieving target equilibria in network routing games without knowing the latency functions. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 31–40, 2014.
- [9] Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer. Stackelberg routing in arbitrary networks. *Math. Oper. Res.*, 35(2):330–346, 2010.
- [10] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 521–530, 2003.
- [11] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.
- [12] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms (3. ed.)*. MIT Press, 2009.
- [13] Stella Dafermos and Anna Nagurney. Sensitivity analysis for the asymmetric network equilibrium problem. *Mathematical programming*, 28(2):174–184, 1984.
- [14] John Fearnley, Martin Gairing, Paul W. Goldberg, and Rahul Savani. Learning equilibria of games via payoff queries. In *ACM Conference on Electronic Commerce*, pages 397–414, 2013.
- [15] Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theor. Comput. Sci.*, 348(2-3):217–225, 2005.
- [16] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multi-commodity networks and generalized congestion games. In *FOCS*, pages 277–285, 2004.
- [17] Steven Fortune, John Hopcroft, and James Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10(2):111–121, 1980.
- [18] Martin Grötschel, László Lovász, and Lex Schrijver. Geometric algorithms and combinatorial optimization. *Algorithms and Combinatorics*, 2:1–362, 1993.

- [19] Tobias Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory Comput. Syst.*, 48(4):781–802, 2011.
- [20] S. Hart and N. Nisan. The query complexity of correlated equilibria. CS arXiv, 2013.
- [21] Albert Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibria in compact games. *Games and Economic Behavior*, 2013. To appear.
- [22] A. Kaporis and P. Spirakis. The price of optimum in Stackelberg games on arbitrary single commodity networks and latency functions. *Theoretical Computer Science*, 410:745–755, 2009.
- [23] George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous users. In *FOCS*, pages 268–276, 2004.
- [24] Elias Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of 16th STACS*, pages 404–413, 1999.
- [25] VS Anil Kumar and Madhav V Marathe. Improved results for Stackelberg scheduling strategies. In *Automata, Languages and Programming*, pages 776–787. Springer, 2002.
- [26] C. H. Papadimitriou. Algorithms, games, and the internet, 2001.
- [27] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.
- [28] J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534, 1965.
- [29] Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Watch and learn: optimizing from revealed preferences feedback. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 949–962, 2016.
- [30] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [31] T. Roughgarden. Routing games. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [32] Tim Roughgarden. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.*, 67(2):341–364, 2003.
- [33] Tim Roughgarden. Stackelberg scheduling strategies. *SIAM J. Comput.*, 33(2):332–350, 2004.
- [34] Tim Roughgarden. On the severity of Braess’s paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.*, 72(5):922–953, 2006.
- [35] Tim Roughgarden and Christos Papadimitriou. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):14, 2008.
- [36] Tim Roughgarden and Florian Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *SODA*, pages 255–267, 2011.
- [37] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- [38] D. Shmoys and C. Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM*, 58:25, 2011.

- [39] Ashish Sureka and Peter R. Wurman. Using Tabu best-response search to find pure strategy Nash equilibria in normal form games. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '05, pages 1023–1029, New York, NY, USA, 2005. ACM.
- [40] Chaitanya Swamy. The effectiveness of Stackelberg strategies and tolls for network congestion games. *ACM Transactions on Algorithms*, 8(4):36, 2012.
- [41] John Glen Wardrop. Some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, volume 1, pages 325–362. Thomas Telford, 1952.
- [42] Michael P. Wellman. Methods for empirical game-theoretic analysis. In *AAAI*, pages 1552–1556, 2006.
- [43] Hai Yang and Hai-Jun Huang. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Research Part B: Methodological*, 38(1):1–15, 2004.
- [44] Hai Yang and Xiaoning Zhang. Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Research Part B: Methodological*, 42(2):99–112, 2008.
- [45] Morteza Zadimoghaddam and Aaron Roth. Efficiently learning from revealed preference. In *Internet and Network Economics - 8th International Workshop, WINE 2012, Liverpool, UK, December 10-12, 2012. Proceedings*, pages 114–127, 2012.

A Proofs from Section 2

Proof of Lemma 2.2. The necessity of the first condition, that $g_e \leq f_e^*$ on every edge, is obvious. For the necessity of the second condition, assume f^* is an equilibrium flow and on some edge $e \notin S$, $g_e < f_e^*$. Then $f_e(g) > 0$ since $f(g) + g = f^*$. By definition of Wardrop equilibrium, then there must exist a path P with $e \in P$ and $l_P(f^*) \leq l_Q(f^*)$ for any path Q . This contradicts that $e \notin S$.

For the sufficiency of the conditions, consider the flow $f^* - g$. This is strictly positive only on shortest-path edges, and hence satisfies the conditions for Wardrop equilibrium with Stackelberg flow g . Since the equilibrium is unique, $f(g) = f^* - g$. ■

B Proofs from Sections 3.1 and 3.2

Proof of Lemma 3.4. Let $(G, l, (s_i, t_i, d_i)_{i \leq k})$ be the given routing game. Recall that we assume that the l_e s satisfy (2)–(4) with $\log K = \text{poly}(\mathcal{I})$. Recall the convex program (1) used to compute the Wardrop equilibrium.

$$\min \Phi(f) := \sum_e \int_0^{f_e} l_e(x) dx \quad \text{s.t.} \quad f = \sum_{i=1}^k f^i, \quad f^i \text{ is an } s_i\text{-}t_i \text{ flow of value } d_i \quad \forall i = 1, \dots, k. \quad (1)$$

Set $\delta = \frac{\epsilon}{4mK}$ and $\varepsilon = \min\left\{\frac{\epsilon(\sum_i d_i)}{2}, \frac{\delta^2}{2K^2}\right\}$. Let \hat{f} be the Wardrop equilibrium, and g be a feasible flow such that $\Phi(g) \leq \Phi(\hat{f}) + \varepsilon$ that we compute in time $\text{poly}(\mathcal{I}, \log(\frac{1}{\varepsilon})) = \text{poly}(\mathcal{I}, \log(\frac{1}{\delta}))$. (We will later require that g is computed via a specific algorithm for solving (1).)

First, we note that given any feasible flow g , one can always obtain an acyclic feasible flow $g' \leq g$ by simply canceling flow along flow-carrying cycles (of each commodity). So in the sequel, we ignore the acyclicity condition and concentrate on obtaining an approximate equilibrium.

Observe that for any feasible flows h, f , we have $\Phi(h) - \Phi(f) \geq v_f \cdot (h - f)$, $v_f = (l_e(f_e))_e$; v_f is called the *subgradient* of Φ at f . So we have

$$\sum_e g_e l_e(\hat{f}_e) - \sum_i d_i D^i(l, \hat{f}) = \sum_e (g_e - \hat{f}_e) l_e(\hat{f}_e) \leq \Phi(g) - \Phi(\hat{f}) \leq \varepsilon.$$

We show below that $\sum_e (g_e - \hat{f}_e) l_e(g_e) \leq \frac{\delta^2}{K}$. Since \hat{f} is an equilibrium, we also have $\sum_e (\hat{f}_e - g_e) l_e(\hat{f}_e) \leq 0$. Adding the two inequalities gives $\sum_e (g_e - \hat{f}_e) (l_e(g_e) - l_e(\hat{f}_e)) \leq \frac{\delta^2}{K}$. Each term in this sum is nonnegative and hence is at most $\frac{\delta^2}{K}$, and therefore we have $|g_e - \hat{f}_e| \leq \delta$ for every edge e (due to inverse- K -continuity). Given this, we have that $l_P(g) \leq l_P(\hat{f}) + mK\delta$ due to the K -Lipschitz condition, and so $D^i(l, g) \leq D^i(l, \hat{f}) + mK\delta$ for every commodity i . Therefore,

$$\begin{aligned} \sum_e g_e l_e(g_e) &\leq \sum_e g_e l_e(\hat{f}_e) + mK\delta \left(\sum_i d_i \right) \leq \sum_i d_i (D^i(l, \hat{f}) + mK\delta) + \varepsilon \\ &\leq \sum_i d_i (D^i(l, g) + 2mK\delta) + \varepsilon \leq \sum_i d_i (D^i(l, g) + \varepsilon). \end{aligned}$$

We now show that $\sum_e (g_e - \hat{f}_e) l_e(g_e) \leq \frac{\delta^2}{K}$. Suppose we obtain the near-optimal solution to (1) by running the ellipsoid method with error parameter $\omega = \frac{\varepsilon}{mK \sum_i d_i}$. This takes time $\text{poly}(\mathcal{I}, \log(\frac{1}{\omega})) = \text{poly}(\mathcal{I}, \log(\frac{1}{\varepsilon}))$. The near-optimality of g then follows from the fact that there exists another feasible flow h satisfying: (i) $\|h - \hat{f}\|_\infty \leq \omega$, and so $\Phi(h) - \Phi(\hat{f}) \leq \sum_e (h_e - \hat{f}_e) l_e(h_e) \leq \omega m (\max_e l_e(h_e)) \leq \omega m K \sum_i d_i = \varepsilon$; (ii) $\sum_e (h_e - g_e) l_e(g_e) = 0$; see, e.g., Sections 3 and 4 and in particular, Lemma 4.5, in [38]. Thus, we have $\sum_e (\hat{f}_e - g_e) l_e(g_e) \geq 0 - \omega m (\max_e l_e(g_e)) \geq -\omega m K \sum_i d_i \geq -\frac{\delta^2}{K}$. ■

Definition of general nonatomic congestion games. This is the following generalization of nonatomic routing games. The edge set is now replaced by a set E of resources, and there are k player-types. Each resource e has a nonnegative, continuous, and strictly increasing delay function, $l_e : \mathbb{R}_+ \mapsto \mathbb{R}_+$. Each player-type i is described by a player-volume d_i and an explicitly-given non-empty strategy set $\mathcal{P}^i \subseteq 2^E$. The combined strategy-choices of the infinitely-many infinitesimal players of each type i can be described by an assignment $f = (f^1, \dots, f^k)$, where $f^i : \mathcal{P}^i \mapsto \mathbb{R}_+$ satisfies $\sum_{P \in \mathcal{P}^i} f_P^i = d_i$; the cost incurred by a strategy $Q \in \bigcup_i \mathcal{P}^i$ is then $l_Q(f) := \sum_{e \in Q} l_e(f_e)$, where $f_e = \sum_{P \in \bigcup_i \mathcal{P}^i: e \in P} f_P^i$. We define $D^i(l, f)$ and an ε -equilibrium as before: so $D^i(l, f) = \min_{P \in \mathcal{P}^i} l_P(f)$, and f is an ε -equilibrium if $\sum_e f_e l_e(f_e) \leq \sum_i d_i (D^i(l, f) + \varepsilon)$. A Nash equilibrium or Nash assignment is a 0-equilibrium, and is known to be unique.

The question with tolls is whether one can impose tolls $\tau \in \mathbb{R}^E$ on resources—which, as before, yield delay functions $(l_e^\tau(x) := l_e(x) + \tau_e)_e$ —in order to achieve a target assignment f^* as the Nash assignment, or ensure that $(f_e^*)_e$ is component-wise close to the Nash assignment.

Proof of Theorem 3.12. We first recall the definition of a Nash equilibrium. A Nash equilibrium of the atomic splittable routing game is a feasible flow f such that $\sum_e f_e^i l_e(f_e) \leq \sum_e g_e^i l_e(f_e - f_e^i + g_e^i)$ for every s_i - t_i flow g^i of value d_i . Equivalently, defining the marginal latency function $\bar{l}_{i,e}(f; x) := l_e(x) + f_e^i l'_e(x)$, where $l'_e(x)$ is the derivative of l_e , this means that if $f_P^i > 0$ for $P \in \mathcal{P}^i$, then P is a shortest s_i - t_i path under the edge costs $(\bar{l}_{i,e}(f; f_e))_e$.

We use the ellipsoid method and dovetail the proof of Theorem 3.2. Given the current ellipsoid center $(\hat{l}, \hat{\tau})$, we obtain a separating hyperplane as in the proof of Theorem 3.2, except that we use the marginal delay functions $(\bar{l}_{i,e}^{\hat{\tau}})_{i,e}$. Let $g = f(\hat{l}^*, \hat{\tau}) = (g^i)_{i \leq k}$ be the flow returned by the oracle. If $g^i = f^{*i}$ for all i , then we are done, so suppose otherwise. Suppose that $f(\hat{l}, \hat{\tau}) \neq f^*$, that is, there is some i such that

$f(\hat{l}, \hat{\tau})^i \neq f^{*i}$. Note that this can be efficiently determined. We can find a player j and paths $P, Q \in \mathcal{P}^j$ such that $f_P^{*j} > 0$, but $\sum_{e \in P} \overline{\hat{l}_{j,e}}(f^*; f_e^*) > \sum_{e \in Q} \overline{\hat{l}_{j,e}}(f^*; f_e^*)$. Thus, the inequality

$$\sum_{e \in P} \overline{l_{j,e}}(f^*; f_e^*) + \tau(P) \leq \sum_{e \in Q} \overline{l_{j,e}}(f^*; f_e^*) + \tau(Q)$$

where both l and τ are variables is violated by $(\hat{l}, \hat{\tau})$ but satisfied by (l^*, τ^*) since (l^*, τ^*) induce f^* (by definition). Notice that the above inequality is indeed linear in l and τ .

Now suppose $f(\hat{l}, \hat{\tau})^i = f^{*i}$ for all i . Then, $g \neq f^*$, we can again find a player j and paths $P, Q \in \mathcal{P}^j$ such that $g_P^j > 0$, but $\sum_{e \in P} \overline{\hat{l}_{j,e}}(g; g_e) > \sum_{e \in Q} \overline{\hat{l}_{j,e}}(g; g_e)$. So consider the inequality

$$\sum_{e \in P} \overline{l_{j,e}}(g; g_e) + \hat{\tau}(P) \leq \sum_{e \in Q} \overline{l_{j,e}}(g; g_e) + \hat{\tau}(Q)$$

where now only the l_e s are variables. This is violated by $(\hat{l}, \hat{\tau})$ but satisfied by (l^*, τ^*) since $g = f(l^*, \tau^*)$. ■

C Proofs from Section 3.3

C.1 Proof of Lemma 3.18

As mentioned in the proof sketch, we first show a property that is weaker than having a discriminating pair. To this end, we define a *good pair of subgraphs* (Definition C.1) and first show in Lemma C.2 that a pair of subgraphs satisfying this weaker property always exist.

Let $(G, \{(s_i, t_i, d_i)\}_{i \in \mathcal{K}})$ be a multicommodity flow instance on a sepa graph. Let \mathcal{H} be the collection of parallel subgraphs of G under a given sepa decomposition tree for G . For any subgraph $H \in \mathcal{H}$ we define the *internal nodes of H* as $V^{\text{int}}(H) := V(H) \setminus \{s_H, t_H\}$. The *internal commodities of H* are $\mathcal{K}^{\text{int}}(H) := \{i \in \mathcal{K} : \{s_i, t_i\} \cap V^{\text{int}}(H) \neq \emptyset\}$. The *external commodities of H* are $\mathcal{K}^{\text{ext}}(H) := \{i \in \mathcal{K} : s_H, t_H \text{ lie on some } s_i\text{-}t_i \text{ path}\}$.

Let $f = (f^i)_{i \in \mathcal{K}}$ and $\tilde{f} = (\tilde{f}^i)_{i \in \mathcal{K}}$ be two feasible multicommodity flows. Define

$$|f_H^i| := \sum_{e=(s_H, v) \in E(H)} f_e^i, \text{ and } |f_H| := \sum_{i \in \mathcal{K}^{\text{ext}}(H)} |f_H^i|.$$

Definition C.1. Given feasible flows f, \tilde{f} in G , subgraphs H_1, H_2 are $(f, \tilde{f}, \mathcal{H})$ -good if:

- (i) the parallel-join of H_1 and H_2 is a subgraph in \mathcal{H} ;
- (ii) $f_e \geq \tilde{f}_e$ for all $e \in E(H_1)$ and $f_e \leq \tilde{f}_e$ for all $e \in E(H_2)$; and
- (iii) $|f_{H_1}| > |\tilde{f}_{H_1}|$ and $|f_{H_2}| < |\tilde{f}_{H_2}|$.

Lemma C.2. For any subgraph H of G , let \mathcal{H}' be the set of subgraphs of H obtained by parallel joins in a given decomposition tree of G , and let f, \tilde{f} be feasible multicommodity flows in G . Either there exists an $(f, \tilde{f}, \mathcal{H}')$ -good pair of subgraphs or one of the following must hold.

1. If $|f_H| = |\tilde{f}_H|$ then $f_e = \tilde{f}_e$ for all $e \in E(H)$.
2. If $|f_H| > |\tilde{f}_H|$ then $f_e \geq \tilde{f}_e$ for all $e \in E(H)$.
3. If $|f_H| < |\tilde{f}_H|$ then $f_e \leq \tilde{f}_e$ for all $e \in E(H)$.

Proof. We proceed by induction on the size of H . In the base case, when H is a single edge, there is no good pair of subgraphs, but one of the three cases clearly holds. For the induction step, suppose H is the join of subgraphs H_1 and H_2 . If H is the parallel join of H_1 and H_2 , then any external commodities of H are external commodities of H_1 and H_2 as well; similarly, any external commodities of H_1 and H_2 are external commodities of H as well. Hence $|f_H| = |f_{H_1}| + |f_{H_2}|$. Note that if $|f_{H_i}| > |\tilde{f}_{H_i}|$ and $|f_{H_j}| < |\tilde{f}_{H_j}|$ for $i \neq j$ and $i, j \in \{1, 2\}$ then H_1 and H_2 form a good pair.

To verify the three cases, suppose $|f_H| = |\tilde{f}_H|$. If $|f_{H_i}| = |\tilde{f}_{H_i}|$ for $i \in \{1, 2\}$, then by induction $f_e = \tilde{f}_e$ for $e \in E(H)$. Otherwise, by the induction hypothesis for $i \neq j$ and $i, j \in \{1, 2\}$, $|f_{H_i}| > |\tilde{f}_{H_i}|$ and $|f_{H_j}| < |\tilde{f}_{H_j}|$ yielding a good pair. If $|f_H| > |\tilde{f}_H|$ then again, either $|f_{H_i}| > |\tilde{f}_{H_i}|$ and $|f_{H_j}| < |\tilde{f}_{H_j}|$ yielding a good pair, or $|f_{H_i}| > |\tilde{f}_{H_i}|$ and $|f_{H_j}| = |\tilde{f}_{H_j}|$. In this case, by induction, $f_e \geq \tilde{f}_e$ for all $e \in E(H)$.

Now suppose H_1 and H_2 are in series. In this case, note that $\mathcal{K}^{\text{ext}}(H) \subseteq \mathcal{K}^{\text{ext}}(H_i)$ for $i \in \{1, 2\}$. Further, if commodity $i \in \mathcal{K}^{\text{ext}}(H_1)$ but $i \notin \mathcal{K}^{\text{ext}}(H)$, then t_i must be an internal node of H_2 . Since every s_i - t_i path contains s_{H_1} , and f, \tilde{f} are feasible flows in G , $|f_{H_1}^i| = |\tilde{f}_{H_1}^i|$. Similarly, if commodity $i \in \mathcal{K}^{\text{ext}}(H_2)$ but $i \notin \mathcal{K}^{\text{ext}}(H)$, then s_i must be an internal node of H_1 . Since every s_i - t_i path contains s_{H_2} , $|f_{H_2}^i| = |\tilde{f}_{H_2}^i|$. Thus,

$$\begin{aligned} |f_{H_1}| - |\tilde{f}_{H_1}| &= \sum_{i \in \mathcal{K}^{\text{ext}}(H_1) \cap \mathcal{K}^{\text{ext}}(H)} \left(|f_{H_1}^i| - |\tilde{f}_{H_1}^i| \right) + \sum_{i \in \mathcal{K}^{\text{ext}}(H_1) \setminus \mathcal{K}^{\text{ext}}(H)} \left(|f_{H_1}^i| - |\tilde{f}_{H_1}^i| \right) \\ &= \sum_{i \in \mathcal{K}^{\text{ext}}(H_1) \cap \mathcal{K}^{\text{ext}}(H)} \left(|f_{H_1}^i| - |\tilde{f}_{H_1}^i| \right) \\ &= \sum_{i \in \mathcal{K}^{\text{ext}}(H)} \left(|f_{H_1}^i| - |\tilde{f}_{H_1}^i| \right) = |f_H| - |\tilde{f}_H| \end{aligned}$$

Similarly, $|f_{H_2}| - |\tilde{f}_{H_2}| = |f_H| - |\tilde{f}_H|$. By induction, either there is a good subgraph, or one of the three cases in the lemma must hold. \blacksquare

Proof of Lemma 3.18. Since f and \tilde{f} are feasible multicommodity flows and $f \neq \tilde{f}$, Lemma C.2 implies that there is an $(f, \tilde{f}, \mathcal{H})$ -good pair of subgraphs H_1, H_2 . So (a) $f_e \geq \tilde{f}_e$ for all $e \in E(H_1)$ and $f_e \leq \tilde{f}_e$ for all $e \in E(H_2)$, and (b) $|f_{H_1}| > |\tilde{f}_{H_1}|$ and $|f_{H_2}| < |\tilde{f}_{H_2}|$. If $f_e > \tilde{f}_e$ for all $e \in E(H_1)$, then we are done. So suppose otherwise.

In the fixed decomposition tree of G , consider the subgraphs in the subtree rooted at subgraph H_1 . Let K be a minimal subgraph that contains both $f_e > \tilde{f}_e$ edges and $f_e = \tilde{f}_e$ edges; that is, every subgraph of K only contains $f_e > \tilde{f}_e$ edges or $f_e = \tilde{f}_e$ edges but not both. Let K be the join of K_1 and K_2 , where K_1 contains $f_e > \tilde{f}_e$ edges. If K_1, K_2 are in parallel, then K_1, K_2 is an $(f, \tilde{f}, \mathcal{H})$ -discriminating pair.

To complete the proof, we show that it cannot be that K_1 and K_2 are in series. Let v be the node joining K_1 and K_2 , so all edges with v at their head lie in $E(K_1)$, and all edges with v at their tail lie in $E(K_2)$. Given a feasible multicommodity flow h , define $b_v(h) = \sum_{(v,u) \in E} h_{v,u} - \sum_{(u,v) \in E} h_{u,v}$. Observe that $b_v(h)$ is simply the node balance $\sum_{i:v=s_i} d_i - \sum_{i:v=t_i} d_i$, and is thus independent of the multicommodity flow. Therefore, $b_v(f) = b_v(\tilde{f})$. Rearranging, this gives $\sum_{e \in E(K_1):e=(v,u)} (f_e - \tilde{f}_e) = \sum_{e \in E(K_2):e=(v,u)} (f_e - \tilde{f}_e)$, which is a contradiction. \blacksquare

D Proofs from Section 3.4

Proof of Claim 3.23. We assume that in τ , there is a single edge $e' = (u, w)$ with negative tolls. If there are multiple such edges, simply repeating the procedure in this proof gives the required tolls τ' . If $f_{e'}(\tau) = 0$, increasing the toll on this edge does no change the equilibrium flow. Hence we assume that $f_{e'}(l, \tau) > 0$.

Let E^+ be the edge set of the graph if it is acyclic; otherwise, let E^+ be the set of edges with strictly positive flow in $g = f(l, \tau)$. Since g is an equilibrium flow, the set of edges E^+ is acyclic. Let $\sigma(v)$ be an ordering of the vertices given by a topological sort of the graph (V, E^+) . Define $S = \{v \in V : \sigma(v) \leq \sigma(u)\}$, where $e' = (u, v)$ is the edge with negative toll. Then $s \in S$ and $t \in V \setminus S$. Let τ' be the tolls obtained by adding $-\tau_{e'}$ to every edge $e \notin E^+$, and also to every edge $e = (x, y) \in E^+$ across the cut $(S, V \setminus S)$. That is,

$$\tau'_{xy} = \begin{cases} \tau_{xy} - \tau_{e'} & \text{if } x \in S, y \in V \setminus S, \text{ or } (x, y) \notin E^+ \\ \tau_{xy} & \text{otherwise.} \end{cases}$$

By this procedure, the toll does not decrease on any edge and increases to zero on edge e' . We claim that the flow at equilibrium remains unchanged. Consider first a path P with $g_P > 0$. All edges $e \in P$ are in E^+ , and exactly one edge crosses the cut $(S, V \setminus S)$. Hence the delay on every such path increases by exactly $-\tau_{e'}$. On any other path, there is at least one edge $e \notin E^+$, hence the delay these paths increases by at least $-\tau_{e'}$. The flow g is thus a flow on shortest paths with tolls τ' , and hence $g = f(l, \tau')$. ■

Proof of Theorem 3.24. We first prove (iii). Let $\tau := \mathbf{1}_r \delta$. Let Φ be the potential function as defined in (1) for the delay functions in Γ , and Φ^τ be the potential function with delay functions that include the toll τ . Note that for any flow f , $\Phi^\tau(f) = \Phi(f) + \tau_r f_r$. Suppose for a contradiction that $f_r(\tau) > f_r(0)$. Then

$$\Phi^\tau(f(0)) = \Phi(f(0)) + \tau_r f_r(0) < \Phi(f(0)) + \tau_r f_r(\tau) < \Phi(f(\tau)) + \tau_r f_r(\tau) = \Phi^\tau(f(\tau)).$$

But this is a contradiction, since $f(\tau)$ is the unique minimizer of Φ^τ .

We now prove part (iv) of the theorem. Let $\tau := -\mathbf{1}_r \delta$. We first prove the lemma for the case that $\mathcal{S}(l, f(\tau)) = \mathcal{S}(l, f(0))$, and then extend it to the case when the set of shortest-path edges differ. For two flows f and g of the same value in G , the difference $h = f - g$ is a circulation and is possibly negative on some edges. If $h_{uv} > 0$ then (u, v) is a forward edge, and if $h_{uv} < 0$ then (u, v) is a backward. We use E^+ and E^- for the set of forward and backward edges respectively.

We want to define a decomposition of h along cycles. For this, let D be the directed graph with the same vertex set as G , but with $(u, v) \in E(D)$ if $(u, v) \in E$ and $h_{uv} > 0$, and $(v, u) \in E(D)$ if $(u, v) \in E$ and $h_{uv} < 0$. Then h defines a circulation \tilde{h} in graph D , where $\tilde{h}_{uv} = h_{uv}$ if (u, v) is a forward edge, and $\tilde{h}_{vu} = -h_{uv}$ if (u, v) is a backward edge. Let $\{\tilde{h}_C\}_{C \in \mathcal{C}}$ be a decomposition of \tilde{h} along directed cycles in D . Then for $(u, v) \in E^+$, $h_{uv} = \sum_{C: (u,v) \in C} \tilde{h}_C$, and for $(u, v) \in E^-$, $h_{uv} = -\sum_{C: (v,u) \in C} \tilde{h}_C$.

Let edge $r = (x, y)$. We will show that (y, x) is in every cycle C . For a contradiction, suppose there exists $C' \in \mathcal{C}$ so that $(y, x) \notin C'$. For any edge $e \in E^+$, $f_e(\tau) > f_e(0)$, and for any edge $e \in E^-$, $f_e(\tau) < f_e(0)$. Further, since $\mathcal{S}(f(0)) = \mathcal{S}(f(\tau))$, the sum of latencies along cycle C' must be zero for both flows $f(\tau)$ and $f(0)$. However,

$$\begin{aligned} \sum_{e \in C'} l_e(f_e(\tau)) &= \sum_{e \in E^+ \cap C'} (l_e(f_e(\tau)) + \tau_e) - \sum_{e \in E^- \cap C'} (l_e(f_e(\tau)) + \tau_e) \\ &> \sum_{e \in E^+ \cap C'} (l_e(f_e(0)) + \tau_e) - \sum_{e \in E^- \cap C'} (l_e(f_e(0)) + \tau_e) \\ &\geq \sum_{e \in E^+ \cap C'} l_e(f_e(0)) - \sum_{e \in E^- \cap C'} l_e(f_e(0)) = 0. \end{aligned}$$

where the second inequality is because $\tau_e = 0$ for $e \neq r$, and $r \notin E^- \cap C'$. This is a contradiction, since the sum of latencies along cycle C' must be zero for flow $f(\tau)$. Thus, for every cycle $C \in \mathcal{C}$, (y, x) must be in C .

Then

$$f_{uv}(\tau) = f_{uv}(0) + \sum_{C \in \mathcal{C}: (u,v) \in C} f_C - \sum_{C \in \mathcal{C}: (v,u) \in C} f_C.$$

Since by the claim edge r is a backward edge in every cycle $C \in \mathcal{C}$, $|f_r(\tau) - f_r(0)| = |\sum_{C \in \mathcal{C}} f_C|$, which is obviously an upper bound on the change in flow on any edge.

We now extend the lemma for the case where $\mathcal{S}(f(0)) \neq \mathcal{S}(f(\tau))$. In fact, we show that for any $\epsilon > 0$, $|f_r(\tau) - f_r(0)| \geq \|f(\tau) - f(0)\|_\infty - \epsilon$. Pick $\nu = \epsilon^2 / (Kd2^{2m})$, where m is the number of edges. Let $a_0 = 0$. For any a_i we define

$$b_i = \sup\{x \in [a_i, \delta] : \mathcal{S}(f(-\mathbb{1}_r x)) = \mathcal{S}(f(-\mathbb{1}_r a_i))\}.$$

and $a_{i+1} = b_i + \nu$. Let j be such that $\delta \in [a_j, a_{j+1}]$. By definition, either $\delta = b_j$ or $\delta \in [b_j, a_{j+1}]$. Since the number of possible sets of shortest-path edges is 2^m , $j \leq 2^m$. Also, for all i , by the first part of the lemma and by continuity of equilibrium flow, $|f_r(-\mathbb{1}_r a_i) - f_r(-\mathbb{1}_r b_i)| \geq \|f(-\mathbb{1}_r a_i) - f(-\mathbb{1}_r b_i)\|_\infty$. Further by Corollary 3.25, $\|f(-\mathbb{1}_r b_i) - f_r(-\mathbb{1}_r a_{i+1})\| \leq \sqrt{Kd\nu}$. Summing up,

$$\begin{aligned} \|f(0) - f(-\mathbb{1}_r \delta)\|_\infty &\leq \sum_{i=0}^j \|f(-\mathbb{1}_r a_i) - f(-\mathbb{1}_r a_{i+1})\|_\infty \\ &\leq \sum_{i=0}^j \|f(-\mathbb{1}_r a_i) - f(-\mathbb{1}_r b_i)\|_\infty + \sum_{i=0}^j \|f(-\mathbb{1}_r b_i) - f(-\mathbb{1}_r a_{i+1})\|_\infty \\ &\leq \sum_{i=0}^j |f_r(-\mathbb{1}_r a_i) - f_r(-\mathbb{1}_r b_i)| + 2^m \sqrt{Kd\nu} \\ &\leq |f_r(0) - f_r(-\mathbb{1}_r \delta)| + \epsilon \end{aligned}$$

where the last inequality follows by the monotonicity of f_r as a function of the toll on edge r . By taking limits, $\|f(0) - f(\tau)\|_\infty \leq |f_r(0) - f_r(\tau)|$. \blacksquare

E Proof of Theorem 4.6

In the following discussion, we fix a Stackelberg flow g , and define $\hat{f} = f(g)$. Flow \tilde{f} is the flow returned by the ϵ -oracle. Define $\delta_e = \tilde{f}_e + g_e - f_e^*$, and for subgraph H , $\delta_H = \|\tilde{f}_H + g_H\| - \|f_H^*\|$. By Lemma 3.5, if \tilde{f} is the flow returned by the ϵ -oracle with Stackelberg flow g , then $\|\hat{f} - \tilde{f}\|_\infty \leq \epsilon'$. Further, recall that by our assumptions on the delay functions,

$$\frac{dl_e(x)}{dx} \geq \frac{1}{K}, \text{ and } |l_e(x) - l_e(y)| \leq K|x - y|$$

for all $x, y \in [0, d]$.

Claim E.1. Fix an edge r with $\delta_r \geq 2(m+1)\epsilon'$. Let H be a subgraph with $r \in E(H)$ so that for every subgraph H' of H with $r \in E(H')$, $\delta_{H'} \geq 2(m+1)\epsilon'$. Then there is an s_H - t_H path P with $r \in P$ and $\delta_e \geq 2\epsilon'$ on every edge $e \in P$.

Proof. We first show that for any subgraph H' if $\delta_{H'} \geq 2m\epsilon'$, then there is an $s_{H'}-t_{H'}$ path P' with $\delta_e \geq 2\epsilon'$ on every edge $e \in P'$. Otherwise, removing all edges with $\delta_e < 2\epsilon'$ must disconnect the terminals of H' . However, the total flow on these edges is at most $2m\epsilon'$, while there is net flow of $2(m+1)\epsilon'$ between the terminals. Hence there must be an $s_{H'}-t_{H'}$ path on which every edge has $\delta_e \geq 2\epsilon'$.

We now prove the claim by induction. If H is a single edge, the claim is obviously true. If H is the series-composition of H_1 and H_2 , say $r \in E(H_1)$. Since $\delta_{H_2} = \delta_H \geq 2(m+1)\epsilon'$, by the above proof, there is an $s_{H_2}-t_{H_2}$ path P_2 on which every edge has $\delta_e \geq 2\epsilon'$. Further, by the inductive hypothesis, there is an $s_{H_1}-t_{H_1}$ path P_1 with $r \in P_1$ and on which every edge has $\delta_e \geq 2\epsilon'$. Path P is then the concatenation of P_1 and P_2 . Lastly, if H is the parallel-composition of H_1 and H_2 with $r \in E(H_2)$, such a path must already exist in H_1 by the inductive hypothesis. This is the path P required in the claim. \blacksquare

Lemma E.2. *If subgraphs H_1, H_2 are $(\tilde{f} + g, f^*)$ - ϵ' -good and $g_e \leq f_e^*$ on every edge, then every edge in H_2 is a non-shortest-path edge with edge costs $(l_e^*(f_e^*))_e$.*

Proof. By Lemma 3.5, $\|\hat{f} - \tilde{f}\|_\infty \leq \epsilon'$. Let s', t' be the common terminals of H_1 and H_2 . Let P be the $s'-t'$ path obtained from Claim E.1 with $r \in P$, and Q be an arbitrary $s'-t'$ path with edges in H_2 . Since $\hat{f}_e + g_e - f_e^* \geq 2\epsilon'$ for all edges in P by the claim, $\hat{f}_e + g_e - f_e^* > 0$. Hence

$$l_P(\hat{f} + g) - l_P(f^*) > l_r(\hat{f}_r + g_r) - l_r(f_r^*) \geq \frac{1}{K}(\hat{f}_r + g_r - f_r^*) \geq \frac{(2mK)^{2m'} - 1}{K}\epsilon'. \quad (8)$$

On path Q ,

$$\begin{aligned} l_Q(\hat{f} + g) - l_Q(f^*) &= \sum_{e \in Q} l_e(\hat{f}_e + g_e) - l_e(f_e^*) \leq \sum_{e \in Q} K|\hat{f}_e + g_e - f_e^*| \\ &\leq \sum_{e \in Q} K \left(|\hat{f}_e - \tilde{f}_e| + |\tilde{f}_e + g_e - f_e^*| \right) \leq mK \left(1 + (2mK)^{2|E(H_2)|} \right) \epsilon'. \end{aligned} \quad (9)$$

Since $\hat{f}_e + g_e - f_e^* > 0$ on every edge in P and $g_e \leq f_e^*$, $\hat{f}_P > 0$, and hence $l_P(\hat{f} + g) \leq l_Q(\hat{f} + g)$. Then from (8) and (9),

$$\begin{aligned} l_Q(f^*) - l_P(f^*) &> \frac{(2mK)^{2m'} - 1}{K}\epsilon' - mK \left(1 + (2mK)^{2|E(H_2)|} \right) \epsilon' \\ &\geq (2mK)^{2m'}\epsilon' \left(\frac{1}{K} - \frac{1}{4mK} \right) - \frac{\epsilon'}{K} - mK\epsilon' \\ &\geq (2mK)^2\epsilon' \left(\frac{1}{K} - \frac{1}{4mK} \right) - \frac{\epsilon'}{K} - mK\epsilon' > 0. \end{aligned}$$

Thus, any $s'-t'$ path Q containing edges in H_2 must be a non-shortest-path, and hence every edge in H_2 is a non-shortest-path edge. \blacksquare

Lemma E.3. *Any subgraph H with $\delta_H < 2(m+1)\epsilon'$ and that contains an edge r with $\delta_r \geq (2mK)^{2|E(H)|}\epsilon'$, must contain a $(\tilde{f} + g, f^*)$ - ϵ' -good pair of subgraphs.*

Proof. In the base case, when graph H consists of a single edge e , the lemma is trivially true since the two conditions cannot be satisfied together.

For the inductive step, let H_1 be the largest subgraph so that $r \in E(H_1)$ and $\delta_L \geq 2(m+1)\epsilon'$ for each subgraph L of H_1 with $r \in E(L)$. Then H_1 must be in parallel with subgraph H_2 , and $\delta_{H_2} < 0$. If

$\delta_e \leq (2mK)^{2|E(H_2)|}\epsilon'$ on each edge in H_2 , then H_1 and H_2 form a good pair of subgraphs. Otherwise, $\delta_{H_2} < 0$ and there is an edge r in H_2 with $\delta_r \geq (2mK)^{2|E(H_2)|}\epsilon'$. By the inductive hypothesis, H_2 must contain a good pair of subgraphs. ■

Corollary E.4 follows from the lemma by observing that $\delta_G = 0$, and if $\|\tilde{f} + g - f^*\|_\infty \geq m(2mK)^{2m}\epsilon'$ then there exists an edge with $\delta_e \geq (2mK)^{2m}\epsilon'$.

Corollary E.4. *If $\|\tilde{f} + g - f^*\| \geq m(2mK)^{2m}\epsilon'$ then there exists a $(\tilde{f} + g, f^*)$ - ϵ' -good pair of subgraphs.*

Proof of Theorem 4.6. Following the proof of Theorem 4.4, in every iteration, $|\bar{S}|$ increases by at least 1 since $f_{H_2} + g_{H_2} < f_{H_2}^*$ and g saturates every edge in \bar{S} . When $S = E$, $g = f^*$. So the algorithm terminates in at most m iterations with some flow g so that $\|\tilde{f} + g - f^*\|_\infty \leq (2mK)^{2m}\epsilon'$. To see that $|g| \leq |g^*|$, note that by Lemma E.2, every edge in H_2 must be a non-shortest-path edge, and hence every edge added to \bar{S} is a non-shortest-path edge. ■

F Proof of Theorem 3.21

We will use the following claim frequently in our proofs.

Claim F.1. *Let subgraph H be the parallel composition of H_1 and H_2 , and τ, τ' be tolls so that, for some $\gamma \in \mathbb{R}$,*

$$\tau'_e = \begin{cases} \tau_e + \gamma & \text{if } e = (s_H, w) \in E(H) \\ \tau_e & \text{if } e \neq (s_H, w) \in E(H) \end{cases}$$

Then

$$(|f_{H_1}(\tau)| - |f_{H_1}(\tau')|) \cdot (|f_{H_2}(\tau)| - |f_{H_2}(\tau')|) \geq 0.$$

Note that the claim does not impose any condition on tolls outside H .

Proof. Let $f = f(\tau)$, $f' = f(\tau')$. For a contradiction, assume without loss of generality that $|f_{H_1}| > |f'_{H_1}|$ but $|f_{H_2}| < |f'_{H_2}|$. Then there are s_H - t_H paths P, Q in H_1, H_2 respectively so that for all $e \in P$, $f_e > f'_e$ and for all $e \in Q$, $f_e < f'_e$. Note that since both paths have a single edge incident to s_H , $\tau'(P) = \tau(P) + \gamma$ and $\tau'(Q) = \tau(Q) + \gamma$. Then, since $f_P > 0$, $f'_Q > 0$,

$$l_P^\tau(f) \leq l_Q^\tau(f) < l_Q^{\tau'}(f') = l_Q^{\tau'}(f') - \gamma \leq l_P^{\tau'}(f') - \gamma = l_P^{\tau'}(f').$$

However, $l_P^{\tau'}(f') < l_P^\tau(f)$, giving us a contradiction. ■

The structure of the proof is as follows. Our algorithm proceeds by considering all subgraphs in \mathcal{H} in a bottom-up manner, i.e., we consider all subgraphs of H before we consider H . Within any subgraph H , we show that by a binary search we can obtain tolls τ so that $|f_H(\tau)| = |f_H^*|$ (Claim F.2). Further, if H is the parallel composition of H_1 and H_2 , we show that by varying just two parameters $\gamma, \gamma^1 \in \mathbb{R}$, we can obtain tolls τ so that $|f_{H_1}(\tau)| = |f_{H_1}^*|$ and $|f_{H_2}(\tau)| = |f_{H_2}^*|$ (Lemma F.3). We use this in our algorithm to show that if the oracle was exact, by conducting suitable binary searches in all subgraphs $H \in \mathcal{H}$ in a bottom-up manner we could obtain tolls τ so that $f_e(\tau) = f_e^*$ on every edge.

However, since we have an ϵ -oracle, we show that for any tolls τ, τ' and subgraph H , if τ and τ' are the same within subgraph H , then the difference in the total flow in H $|f_H(\tau) - f_H(\tau')|$ is an upper bound

on the difference $|f_e(\tau) - f_e(\tau')|$ on any edge in $E(H)$ (Lemma F.4). This allows us to show that in our bottom-up approach, as long as in every subgraph we obtain tolls so that we have approximately the correct amount of flow, the flow on every edge is approximately correct as well. Finally, Claim F.6 bounds the tolls after every subgraph is processed, allowing us to obtain bounds on the number of binary searches and thus the runtime.

Claim F.2. *For any tolls $(\tau_e)_{e \in E}$ and subgraph H , there exists $\gamma \in \mathbb{R}$ so that the tolls $(\tau'_e)_{e \in E}$ given by*

$$\tau'_e = \begin{cases} \tau_e + \gamma & \text{if } e = (s_H, w) \\ \tau_e & \text{otherwise} \end{cases}$$

satisfy $|f_H(\tau')| = |f_H^*|$. Further, $\gamma \in [-N(\tau), N(\tau)]$.

Proof. Since τ' is defined by γ , we vary γ in the proof and assume τ' varies accordingly. We will show that $|f_H(\tau')| = 0$ and d for $\gamma = -N(\tau)$ and $N(\tau)$ respectively. Since $0 \leq |f_H^*| \leq d$, by the continuity of flow with respect to tolls (Theorem 3.24(ii)), the claim is true.

For $\gamma = -N(\tau)$, let $f = f(\tau')$, and suppose $f_Q > 0$ for some path $Q \in \mathcal{P}$ that does not contain an edge from H . Let $R \in \mathcal{P}$ be a path with some edge $(s_H, w) \in E(H) \cap R$. Then, since $\tau'(Q) = \tau(Q)$ and latencies are strictly increasing,

$$\begin{aligned} l_Q(f) + \tau'(Q) &> \min_P \tau(P) \\ &= ml_{\max}(d) + \max_P \tau(P) - ml_{\max}(d) - \max_P \tau(P) + \min_P \tau(P) \\ &= ml_{\max}(d) + \max_P \tau(P) - N(\tau) \geq l_R(f) + \tau'(R) \end{aligned}$$

and hence, f cannot be an equilibrium flow with tolls τ' . Thus, $|f_H(\tau')| = d$. A similar argument proves that if $\gamma = N(\tau)$, then $f_e(\tau') = 0$ for any edge $e = (s_H, w) \in E(H)$, and hence $|f_H(\tau')| = 0$. \blacksquare

Lemma F.3. *Let H be the parallel composition of H_1 and H_2 and τ be arbitrary edge tolls. There exist γ^* , $\gamma_1^* \in \mathbb{R}$ so that the tolls τ'' defined by*

$$\tau''_e = \begin{cases} \tau_e + \gamma^* + \gamma_1^* & \text{if } e = (s_H, w) \in E(H_1) \\ \tau_e + \gamma^* & \text{if } e = (s_H, w) \in E(H) \setminus E(H_1) \\ \tau_e & \text{otherwise} \end{cases}$$

satisfy $|f_{H_1}(\tau'')| = |f_{H_1}^*|$ and $|f_{H_2}(\tau'')| = |f_{H_2}^*|$. Further, $\gamma_1^* \in [-2N(\tau), 2N(\tau)]$, and $\gamma^* \in [-N(\tau), N(\tau)]$.

Proof. We first obtain γ_1^* , and then γ^* . To obtain γ_1^* , we want to fix the total flow on H_1 to be $|f_{H_1}^*|$. For this, consider the single-commodity routing game Γ_H restricted to subgraph H with s_H and t_H as terminals for the single commodity, demand $d_H = |f_H^*|$, and tolls τ . Then by Claim F.2 there exists $\gamma_1 \in [-N(\tau), N(\tau)]$ so that, if τ' are the tolls defined in the statement of the claim and $g = f(\tau')$, then $|g_{H_1}| = |f_{H_1}^*|$. Further, in the original routing game Γ with tolls τ' so obtained, by Claim F.2 there exists $\gamma \in [-2N(\tau), 2N(\tau)]$ so that, with tolls τ'' defined as $\tau''_e = \tau'_e + \gamma$ for $e = (s_H, w) \in E(H)$ and $\tau''_e = \tau'_e$ otherwise, $|f_H(\tau'')| = |f_H^*|$. We claim that $|f_{H_1}(\tau'')| = |f_{H_1}^*|$. This is sufficient to prove the lemma.

Let $f = f(\tau'')$. We prove the claim by contradiction. Suppose that $|f_{H_1}| > |f_{H_1}^*|$ (the case where $|f_{H_1}| < |f_{H_1}^*|$ can be dealt with similarly). Then $|f_{H_2}| < |f_{H_2}^*|$. This implies that $|f_{H_1}| > |g_{H_1}|$ and $|f_{H_2}| < |g_{H_2}|$, where $g = f(\tau')$ in the game Γ_H . Thus there exist s_H - t_H paths P and Q in H_1 and H_2 respectively so that for all $e \in P$, $f_e > g_e$ and for all $e \in Q$, $f_e < g_e$. Since f is an equilibrium flow with tolls τ'' and g is an equilibrium flow with tolls τ' , and by construction of τ'' ,

$$l_P''(f) \leq l_Q''(f) < l_Q''(g) = l_Q'(g) + \gamma \leq l_P'(g) + \gamma = l_P''(g).$$

However, $l_P''(g) < l_P''(f)$, giving us a contradiction. ■

Lemma F.4. *Let H be a subgraph of G and τ, τ' be tolls so that, for some $\gamma \in \mathbb{R}$,*

$$\tau'_e = \begin{cases} \tau_e + \gamma & \text{if } e = (s_H, w) \in E(H) \\ \tau_e & \text{if } e \neq (s_H, w) \in E(H) \end{cases}$$

Then for any subgraph K of H ,

$$||f_K(\tau)| - |f_K(\tau')|| \leq ||f_H(\tau)| - |f_H(\tau')||.$$

Proof. Let $f = f(\tau)$ and $f' = f(\tau')$. The proof is by induction on the length of the path between H and K in the decomposition tree. In the base case, $H = K$, and the lemma is obviously true. Otherwise, H is the composition of H_1 and H_2 , and K is a subgraph of either H_1 or H_2 . Assume that K is a subgraph of H_1 , and note that tolls τ, τ' satisfy the condition in the lemma for subgraphs H_1 and H_2 as well.

If H is the series-composition of H_1 and H_2 , by induction,

$$||f_K| - |f'_K|| \leq ||f_{H_1}| - |f'_{H_1}||.$$

Since H consists of H_1 and H_2 in series, $|f_{H_1}| = |f_H|$ and $|f'_{H_1}| = |f'_H|$, hence the lemma is true in this case.

If H is the parallel composition of H_1 and H_2 , then by Claim F.1,

$$(|f_{H_1}| - |f'_{H_1}|) \cdot (|f_{H_2}| - |f'_{H_2}|) \geq 0. \tag{10}$$

Then

$$\begin{aligned} ||f_K| - |f'_K|| &\leq ||f_{H_1}| - |f'_{H_1}|| \leq ||f_{H_1}| - |f'_{H_1}|| + ||f_{H_2}| - |f'_{H_2}|| \\ &= ||f_{H_1}| - |f'_{H_1}| + |f_{H_2}| - |f'_{H_2}|| \\ &= ||f_H| - |f'_H||, \end{aligned}$$

where the first inequality follows by induction, the first equality follows from (10), and the second equality is since H is the parallel composition of H_1 and H_2 . ■

For the proof of Theorem 3.21, we first show the correctness and number of steps required for each subgraph in Step T2 (Claim F.5), bound the increase in tolls in each step (Claim F.6), and show that when Step T2 completes for subgraph $H \in \mathcal{H}$ the flow on every edge of H is approximately equal to the target flow (Claim F.7).

Claim F.5. *For any subgraph $H \in \mathcal{H}$, (i) each iteration of Step T2 requires at most $\log(4N(\tau)/\kappa(\epsilon/m))$ queries to find γ , and (ii) Step T2 completes in at most $\log(2N(\tau)/\kappa(\epsilon/m))$ iterations with tolls τ' that satisfy $||\tilde{f}_{H_i}(\tau')| - |f_{H_i}^*|| \leq 6m\epsilon$ for $i \in \{1, 2\}$.*

Proof. For any $\gamma_1 \in [-N(\tau), N(\tau)]$ by Claim F.2, there exists $\gamma' \in [-2N(\tau), 2N(\tau)]$ so that $|\tilde{f}_H(\tau')| = |f_H^*|$, where tolls τ' are given by

$$\tau'_e = \begin{cases} \tau_e + \gamma' + \gamma_1 & \text{if } e = (s_H, w) \in E(H_1) \\ \tau_e + \gamma' & \text{if } e = (s_H, w) \in E(H) \setminus E(H_1) \\ \tau_e & \text{otherwise.} \end{cases}$$

Further, by Theorem 3.24(i), for any γ that satisfies $|\gamma' - \gamma| \leq \kappa(\epsilon/m)$, the flow on any edge differs from $f(\tau')$ by at most ϵ . Thus γ can be found in at most $\log(4N(\tau)/\kappa(\epsilon/m))$ queries. This completes the proof of the first part of the claim.

Now for a particular iteration of Step T2, let γ, γ_1 result in tolls τ' so that $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$. We will show firstly that in this iteration, the algorithm is correct: if $|\tilde{f}_{H_1}(\tau')| - |f_{H_1}^*| \geq 4m\epsilon$, then $\gamma_1 < \gamma_1^*$. A similar argument shows that if $|\tilde{f}_{H_1}(\tau')| - |f_{H_1}^*| \leq -4m\epsilon$, then $\gamma_1 > \gamma_1^*$. Secondly, we will show that if $|\gamma_1 - \gamma_1^*| \leq \kappa(\epsilon/m)$ and $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$ then $\left| |\tilde{f}_{H_i}(\tau')| - |f_{H_i}^*| \right| \leq 6m\epsilon$ for $i \in \{1, 2\}$, and hence the termination condition in the step must be satisfied in $\log(2N(\tau)/\kappa(\epsilon/m))$ iterations.

For correctness of the binary search procedure, since $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$ and H is the parallel composition of H_1 and H_2 , it follows that

$$|\tilde{f}_{H_2}(\tau')| - |f_{H_2}^*| \leq -2m\epsilon.$$

and hence, since $\|\tilde{f}(\tau') - f(\tau')\|_\infty \leq \epsilon$,

$$|f_{H_2}(\tau')| < |f_{H_2}^*| \text{ and } |f_{H_1}(\tau')| > |f_{H_1}^*|.$$

Let τ^* be the tolls as defined in Lemma F.3. Then there exist s_H - t_H paths P and Q in H_1 and H_2 respectively so that for all $e \in P$, $f_e(\tau') > f_e(\tau^*)$ and for all $e \in Q$, $f_e(\tau') < f_e(\tau^*)$. Then since both flows are equilibrium flows for the respective tolls, and by construction of τ' and τ^* ,

$$\begin{aligned} l_P^{\tau'}(f(\tau')) &\leq l_Q^{\tau'}(f(\tau')) < l_Q^{\tau'}(f(\tau^*)) = l_Q^{\tau^*}(f(\tau^*)) + (\gamma - \gamma^*) \\ &\leq l_P^{\tau^*}(f(\tau^*)) + (\gamma - \gamma^*) = l_P^{\tau^*}(f(\tau^*)) + (\gamma_1^* - \gamma_1). \end{aligned}$$

However, $l_P^{\tau'}(f(\tau^*)) < l_P^{\tau'}(f(\tau'))$, and hence $\gamma_1 < \gamma_1^*$. This proves the correctness of the step.

For termination, we assume as stated earlier that γ, γ_1 are chosen so that $|\gamma_1 - \gamma_1^*| \leq \kappa(\epsilon/m)$ and $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$. Let $\hat{\tau}$ be defined as

$$\hat{\tau}_e = \begin{cases} \tau_e + \gamma + \gamma_1^* & \text{if } e = (s_{H_1}, w) \in E(H_1) \\ \tau_e + \gamma & \text{if } e = (s_H, w) \in E(H) \setminus E(H_1) \\ \tau_e & \text{otherwise} \end{cases}$$

Since $|\gamma_1 - \gamma_1^*| \leq \kappa(\epsilon/m)$, by Theorem 3.24(i),

$$\left| |f_H(\tau')| - |f_H(\hat{\tau})| \right| \leq m\epsilon. \quad (11)$$

Since $\left| |\tilde{f}_H(\tau')| - |f_H^*| \right| \leq 2m\epsilon$, $\left| |f_H(\tau')| - |f_H^*| \right| \leq 3m\epsilon$. Since $|f_{H_i}^*| = |f_{H_i}(\tau^*)|$ for $i \in \{1, 2\}$, by Lemma F.4 and (11),

$$\left| |f_{H_i}(\hat{\tau})| - |f_{H_i}(\tau^*)| \right| \leq \left| |f_H(\hat{\tau})| - |f_H(\tau^*)| \right| \leq \left| |f_H(\tau')| - |f_H(\tau^*)| \right| + m\epsilon \leq 4m\epsilon.$$

Since (11) also holds for H_1 and H_2 , putting these together, we get

$$||f_{H_i}(\tau')| - |f_{H_i}(\tau^*)|| \leq 5m\epsilon.$$

and hence, $||\tilde{f}_{H_i}(\tau')| - |f_{H_i}^*|| \leq 6m\epsilon.$ ■

Claim F.6. *Let H' be the i th subgraph processed by the algorithm, and let τ' be the tolls immediately before H' is processed. Then $N(\tau') \leq 5^i ml_{\max}(d)$.*

Proof. If H' is the first subgraph to be processed, since initially the tolls are zero, the claim holds. Otherwise, let H be the subgraph processed immediately before H' and let τ be the tolls before H is processed. By the inductive hypothesis, $N(\tau) \leq 5^{i-1} ml_{\max}(d)$.

By the algorithm, when H is processed, only the tolls on edges incident to s_H are changed. Since the graph is acyclic, this means that the toll on any path changes on at most a single edge. Further, the difference in tolls on any two edges changes by at most $|\gamma + \gamma_1| \leq |4N(\tau)|$. Hence, the difference in tolls on any two paths changes by at most $|4N(\tau)|$. Thus,

$$\begin{aligned} N(\tau') &= \max_P \tau'_P - \min_P \tau'_P + ml_{\max}(d) \\ &\leq |\max_P \tau'_P - \min_P \tau'_P - (\max_P \tau_P - \min_P \tau_P)| + N(\tau) \\ &\leq 5N(\tau) \leq 5^i ml_{\max}(d). \end{aligned}$$
■

Let K be a subgraph of H . We define the *distance* between H and K as the length of the path connecting the corresponding vertices in the decomposition tree. For any subgraph H of G , let τ^H be the tolls immediately after H is processed by the algorithm.

Note that since the algorithm proceeds in a bottom-up manner, for any subgraph K of H , there exists γ so that $\tau_e^H = \tau_e^K$ on any edge $e \neq (s_K, v) \in E(K)$ and $\tau_e^H = \tau_e^K + \gamma$ on any edge $e = (s_K, v) \in E(K)$.

Claim F.7. *If K is a subgraph of H at a distance i , and H is the parallel composition of two subgraphs, then $||f_K(\tau^H)| - |f_K^*|| \leq 10m(i+1)\epsilon$.*

Proof. The proof is by induction on the distance i . If $i = 0$, i.e. $K = H$, then the proof follows immediately from the termination condition in the algorithm, noting that $|f_H(\tau^H) - \tilde{f}_H(\tau^H)| \leq m\epsilon$. For the inductive case, let \hat{K} and K' be the parent and sibling of K respectively. By the inductive hypothesis, $|f_{\hat{K}}(\tau^H)| - |f_{\hat{K}}^*| \leq 14mi\epsilon$. If \hat{K} is the series composition of K and K' , then since in any flow the flow value in K , K' and \hat{K} must be the same, the proof follows.

If \hat{K} is the parallel composition of K and K' , then from the algorithm,

$$|f_{\hat{K}}(\tau^{\hat{K}}) - f_{\hat{K}}^*| \leq 3m\epsilon. \tag{12}$$

From (12) and the induction hypothesis,

$$|f_{\hat{K}}(\tau^{\hat{K}}) - f_{\hat{K}}(\tau^H)| \leq (10i + 3)m\epsilon. \tag{13}$$

By Lemma F.4, we know that

$$|f_K(\tau^H) - f_K(\tau^{\hat{K}})| \leq |f_{\hat{K}}(\tau^H) - f_{\hat{K}}(\tau^{\hat{K}})| \leq (10i + 3)m\epsilon,$$

where the second inequality follows from (13). But from the algorithm, since \hat{K} is the parallel composition of K and K' ,

$$|f_K(\tau^{\hat{K}}) - f_K^*| \leq 6m\epsilon.$$

The proof immediately follows. ■

Proof of Theorem 3.21. Let $H \in \mathcal{H}$ be the i th subgraph processed by the algorithm, and let τ and τ' be the tolls immediately before and after subgraph H is processed by the algorithm. Then by Claim F.6, $N(\tau) \leq 5^{i-1}ml_{\max}(d)$. By Claim F.5, the algorithm completes processing subgraph H in at most $\log(4N(\tau)/\kappa(\epsilon/m))$ iterations. Since each iteration involves querying the oracle at most $\log(4N(\tau)/\kappa(\epsilon/m))$ times, this implies that the algorithm makes at most

$$\log^2(4m5^{i-1}l_{\max}(d)/\kappa(\epsilon/m)) \leq m^2 \log^2(20ml_{\max}(d)/\kappa(\epsilon/m))$$

queries. Since there are at most m subgraphs, the bound in the statement of the theorem holds. Further, let H be the last subgraph in \mathcal{H} processed by the algorithm. If $H = G$, then by Claim F.7, the second part of the theorem holds immediately. If not, then G is the series composition of a number of subgraphs H_1, H_2, \dots, H_t , each of which is in \mathcal{H} . Then $|f_{H_i}(\tau')| = d$ for each subgraph H_i . Further, if τ^i were the tolls obtained after subgraph H_i was processed by the algorithm, and they were processed in this order (thus $\tau' = \tau^t$), then for each $e \in E(H_i)$ and $j \geq i$, $\tau_e^j = \tau_e^i$. Thus, by Lemma F.4, for any subgraph H_i and any edge $e \in E(H_i)$, $f_e(\tau') = f_e(\tau^i)$. By Claim F.7, $|f_e(\tau^i) - f_e^*| \leq 14m^2\epsilon$. This is true for every edge, completing the proof of the theorem. ■