http://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–35
ISSN 0364-765X (print), ISSN 1526-5471 (online)

Parametric Computation of Minimum-Cost Flows with Piecewise Quadratic Costs

Max Klimm,[a] **Philipp Warode**[a]

[a] Institute of Mathematics, Technische Universität Berlin, 10623 Berlin, Germany
Contact: klimm@math.tu-berlin.de, https://orcid.org/0000-0002-9061-2267 (MK); warode@math.tu-berlin.de, https://orcid.org/0000-0002-2878-6872 (PW)

Recevied: November 18, 2019
Revised: September 10, 2020
Accepted: December 17, 2020
Published Online in Articles in Advance: September 14, 2021

MSC2000 Subject Classification: Primary: 90C31; secondary: 90C20, 90C49, 68Q25, 91A07
OR/MS Subject Classification: Primary: Networks/graphs: flow algorithms; secondary: mathematics: piecewise linear

https://doi.org/10.1287/moor.2021.1151

Abstract. We develop algorithms solving parametric flow problems with separable, continuous, piecewise quadratic, and strictly convex cost functions. The parameter to be considered is a common multiplier on the demand of all nodes. Our algorithms compute a family of flows that are each feasible for the respective demand and minimize the costs among the feasible flows for that demand. For single commodity networks with homogeneous cost functions, our algorithm requires one matrix multiplication for the initialization, a rank 1 update for each nondegenerate step and the solution of a convex quadratic program for each degenerate step. For nonhomogeneous cost functions, the initialization requires the solution of a convex quadratic program instead. For multi-commodity networks, both the initialization and every step of the algorithm require the solution of a convex program. As each step is mirrored by a breakpoint in the output this yields output-polynomial algorithms in every case.

Funding: This work was supported by Deutsche Forschungsgemeinschaft [CRC/TRR 154, Project A07, EXC-2046/1, Project ID: 390685689].

Keywords: parametric quadratic program • parametric flow • computational complexity • graph Laplacian • Wardrop equilibrium

1. Introduction

This paper studies parametric minimum cost flow problems with separable, continuous, piecewise quadratic, and strictly convex costs. Formally, we are given a directed graph $G = (V, E)$ with n vertices and m edges. Every edge e has a cost function $F_e : \mathbb{R} \to \mathbb{R}$ of the form $F_e(x) = \int_0^x f_e(s)\,\mathrm{d}s$ where f_e is a strictly increasing and piecewise linear function. Furthermore, let $\mathbf{q} \in \mathbb{R}^n$ be an arbitrary demand vector with $\sum_{v\in V} q_v = 0$. We are interested in computing a function $\mathbf{x} : \mathbb{R}_{\geq 0} \to \mathbb{R}^m_{\geq 0}$ such that for all $\lambda \geq 0$, the vector $\mathbf{x}(\lambda) = (x_e(\lambda))_{e\in E}$ is a flow that minimizes the separable cost function $\sum_{e\in E} F_e(x_e)$ and is feasible for the balance vector $\lambda\mathbf{q}$, that is, for all $\lambda \geq 0$ the flow $\mathbf{x}(\lambda)$ solves the optimization problem

$$\begin{aligned} \min \quad & \sum_{e\in E} F_e(x_e) \\ \text{s.t.} \quad & \sum_{e\in\delta^-(v)} x_e - \sum_{e\in\delta^+(v)} x_e = \lambda q_v \quad \text{for all } v \in V, \\ & x_e \geq 0 \qquad\qquad\quad \text{for all } e \in E. \end{aligned} \tag{1}$$

On our way to solve the parametric optimization problem (1), we also solve a relaxation of the problem where the non-negativity constraints on the flow are omitted.

The nonparametric variant of these problems has been studied in a variety of contexts. For the special case that all marginal cost functions f_e are linear (and the non-negativity constraints are relaxed), the flow (1) describes an electric network with constant resistances on the edges and constant inflows and outflows given by the demand vector. The flow on the edges corresponds to the electric current, and the slope of the marginal cost functions f_e corresponds to the resistance of the edge. Piecewise linear (and possibly noncontinuous) marginal cost functions as we consider them in this paper allow to model more elaborated electric networks with special components such as diodes or transistors or active components such as batteries or generators on the edges. For nonlinear marginal cost functions f_e and relaxed non-negativity constraints, the system (1) models other physical networks such as natural gas networks or water networks (Gross et al. [24]). Although we are only able to handle piecewise linear marginal cost functions f_e and thus piecewise quadratic cost functions F_e, it is worth noting that any

continuous convex cost function F_e can be approximated with arbitrary precision by a piecewise quadratic continuous convex cost function $\tilde{F}_e$ (Rockafellar [50, section 8]).

We also study a generalization of the parametric problem (1) to multi-commodity flows. Here, we are given a finite set $K = \{1, \dots, k\}$ of commodities and for each commodity $j \in K$ a demand vector $\mathbf{q}_j = (q_{v,j})_{v \in V} \in \mathbb{R}^n$ with $\sum_{v \in V} q_{v,j} = 0$. We are then interested in computing a function $\mathbf{x} : \mathbb{R}_{\geq 0} \to \mathbb{R}^{mk}_{\geq 0}$ such that for all $\lambda \geq 0$, the vector $\mathbf{x}(\lambda) = (x_{e,j}(\lambda))_{e \in E, j \in K}$ is a multi-commodity flow that minimizes the cost function $\sum_{e \in E} F_e(\sum_{j \in K} x_{e,j})$ and for every commodity $j \in K$ the flow $\mathbf{x}_j = (x_{e,j})_{e \in E}$ is feasible for the balance vector $\lambda \mathbf{q}_j$, that is, for all $\lambda \geq 0$, the flow $\mathbf{x}(\lambda)$ solves the optimization problem

$$\begin{aligned} \min \quad & \sum_{e \in E} F_e\left(\sum_{j \in K} x_{e,j}\right) & & \\ \text{s.t.} \quad & \sum_{e \in \delta^-(v)} x_{e,j} - \sum_{e \in \delta^+(v)} x_{e,j} = \lambda q_{v,j} & \text{for all } v \in V,\ j \in K, & \\ & x_{e,j} \geq 0 & \text{for all } e \in E,\ j \in K. & \end{aligned} \tag{2}$$

An important application for the nonparametric variant of the multi-commodity setting are traffic and telecommunication networks. Here, every commodity j corresponds to a population of infinitesimally small users that wish to travel (or send a message) from a source node $s_j \in V$ to a target node $t_j \in V$ in the network. That is, the demand vector of commodity j has the form $q_{v,j} = -1$ if $v = s_j$, $q_{v,j} = 1$ if $v = t_j$, and $q_{v,j} = 0$, otherwise. In this way, all (s_j, t_j) flows are feasible for commodity j. In this setting, the marginal cost functions f_e are interpreted as travel times. When the total flow on edge e is equal to $z = \sum_{j \in K} x_{e,j}$, then every flow particle traversing e experiences a latency of $f_e(z)$. Wardrop [57] postulated an equilibrium principle for equilibrium flows that requires that all flow particles only use paths of minimum travel time. A multi-commodity flow $\mathbf{x}$ is called a *Wardrop equilibrium* if $\sum_{e \in P} f_e(\sum_{j \in K} x_{e,j}) \leq \sum_{e \in Q} f_e(\sum_{j \in K} x_{e,j})$ for all commodities $j \in K$ and (s_j, t_j) paths P, Q with $x_{e,j} > 0$ for all $e \in P$. Beckman et al. [6] show that for continuous marginal cost functions f_e, a multi-commodity flow for demands $\lambda \mathbf{q}$ is a Wardrop equilibrium if and only if it solves (2). Thus, being able to solve (2) allows to compute Wardrop equilibria. In addition, solving (2) for slightly altered marginal cost functions of the form $f^*_e(x) = f_e'(x)x + f_e(x)$ allows to compute system optimal traffic flows that minimize the total travel time.

The nonparametric variants of the single-commodity and multi-commodity mincost flow problems (1) and (2) are well studied in the literature, and many algorithms for their solution have been proposed. For the special case that the cost functions F_e are quadratic (rather than piecewise quadratic), both the nonparametric variant of (1) and (2) can be solved with general quadratic programming techniques such as Frank-Wolfe decomposition (Frank and Wolfe [18]), Lemke's algorithm (Sacher [52]), or the Ellipsoid method (Kozlov et al. [38]), the latter method being polynomial. For the nonparametric variant of the single-commodity problem (1), there exists a weakly polynomial algorithm by Minoux [43, 44] and a strongly polynomial algorithm by Végh [56].

Yet, all the previous algorithms solve the nonparametric variants of (1) or (2) only. For the management of electrical, gas, water, telecommunication, and traffic networks, it is desirable to be able to also solve the parametric version of the problems. This is, in part, because these networks may behave quite paradoxically. Braess' paradox (Braess [9]) shows that in a single-commodity traffic network, the traffic along an edge may increase or decrease when the demand increases. Similarly, in gas networks, the maximal pressure on a node (corresponding to the dual multiplier of the corresponding balance constraint) may increase when decreasing $\mathbf{q}$ (Klimm et al. [36]). In fact, for traffic networks, there is profound interest in understanding the performance of road networks as the demand of all commodities changes by a common factor (O'Hare et al. [49], Youn et al. [58]). Thus far, the computation of the respective Wardrop equilibria and system optimal flows relies on repeated runs of the Frank-Wolfe method for small discretization of the demand space. Our method allows to compute all Wardrop equilibria and system optimal flows for piecewise linear latencies directly and without any discretization error.

1.1. Our Results

We develop algorithms that compute the parametric version of the single-commodity (1) and multi-commodity (2) piecewise quadratic mincost flow problems, respectively. We first explain our results for the parametric single-commodity flow problem (1).

Our algorithm is based on a relaxation of (1) where the non-negativity constraints on the flows are dropped, which we call *electrical flow*. For electrical flows, the dual variables for the balance condition of each node can be interpreted as vertex potentials. We first consider the special case that all marginal cost functions are continuous and homogeneous (i.e., $f_e(0) = 0$). This case corresponds to electric networks with piecewise constant resistances

and only passive components on the edges. In this case, there is piecewise linear bijection between vertex potentials and the demand vector satisfied by the corresponding optimal flows; that is, there is a partition of the space of all vertex potentials into different regions such that within each region the bijection is linear. Our algorithm for the parametric computation of flows starts at the all-zero potential vector that induces the all-zero flow. It computes the corresponding bijection (expressed via the inverse of the Laplacian in the corresponding linear regime) and increases the flow linearly until the region is left that happens when one or more edges enter or leave the support. When the instance is nondegenerate, no two edges arrive at a breakpoint at the same time, and the bijection in the new region can be recomputed from the old one by a simple and computationally inexpensive Simplex-like pivot step. For the degenerate case where two or more edges are at a breakpoint at the same time, we develop a lexicographic rule similar to the lexicographic rule for the Simplex algorithm that allows to navigate to the next region, even in the degenerate case. Alternatively, we can recompute the new bijection with a convex quadratic program.

We then work toward achieving the same guarantees for the original problem with non-negativity constraints for the flows. Actually, we solve a more general problem where the marginal cost functions f_e are allowed to be discontinuous and may also take the values $+\infty$ and $-\infty$. The values $+\infty$ and $-\infty$ can effectively be used in order to impose lower and upper capacities on the flow on the edges; that is, for each edge, there is a lower capacity $l_e \in \mathbb{R} \cup \{-\infty\}$ and an upper capacity $u_e \in \mathbb{R} \cup \{+\infty\}$ and a flow is only feasible when $l_e \le x_e \le u_e$ for all $e \in E$. In particular, by imposing a lower capacity $l_e = 0$, directed flows can be enforced thus solving the parametric version of the directed piecewise quadratic mincost flow problem (1). Being able to handle discontinuous marginal cost functions is also beneficial to handle specific components of electric or gas networks such as diodes or backpressure valves that only let specific flows pass. For traffic networks, discontinuous marginal cost functions allow to compute system optimal flows when the underlying travel time functions are piecewise linear. From a mathematical point of view, discontinuous marginal cost functions are challenging as there is no bijection between demand vectors and vertex potential anymore. We circumvent this issue by introducing a repair step that increases the vertex potential of a connected component of the graph whenever this situation occurs. This repair step allows to transfer all results obtained for continuous marginal cost functions also for the discontinuous case.

Our algorithm is output-polynomial. Specifically, it runs in $\mathcal{O}(n^{2.375} + n^2 L)$ for nondegenerate instances where L is the number of breakpoints of the output flow functions. The $\mathcal{O}(n^{2.375})$ time stems from the need to compute one matrix inverse at the start of the algorithm (for which we use fast matrix multiplication) and the $\mathcal{O}(n^2 L)$ time is because of the need to perform one rank 1 update for each transition between regions. For degenerate instances, instead of performing a series of rank 1 updates per transition, we recompute the bijection after the transition by a convex program and obtain an algorithm that runs in time $\mathcal{O}(n^{2.375} + \mathrm{poly}(n)L)$ (see Table 1 for an overview of the achieved runtimes in the different cases).

Finally, we relax the assumption that the marginal cost functions are homogeneous; that is, we allow marginal cost functions with $f_e(0) \neq 0$. This case occurs, for example, in electrical networks with active elements on the edges or in mincost flow networks with lower capacities on the edges. The main difficulty is that in this case the all-zero flow may not be feasible or the all-zero vector may not induce the all-zero flow. Thus, we introduce an additional phase I during which we compute the minimal λ such that the demand $\lambda \mathbf{q}$ is feasible, the corresponding flow, and the corresponding vertex potentials. This can be achieved by solving a convex quadratic program that can be done in polynomial time by the Ellipsoid method.

For multi-commodity networks, we use a similar approach, but recompute the linear map between vertex potentials and demand vectors by a convex program after entering a new region.

All our results for single-commodity networks translate to a more general problem where we are given a nondecreasing piecewise linear function $\mathbf{h} : \mathbb{R}_{\ge 0} \to \mathbb{R}^n$ with $\sum_{v \in V} h_v(\lambda) = 0$ for all $\lambda \ge 0$ and we wish to compute a family of flows $\mathbf{x} : \mathbb{R}_{\ge 0} \to \mathbb{R}^m$ such that $\mathbf{x}(\lambda)$ is a mincost flow for the demand $\mathbf{h}(\mathbf{x})$. Given a starting solution, we can still use our mechanisms to increase $\mathbf{h}(\lambda)$ linearly. Because we need to compute a start solution first, we get

Table 1. Running time of our algorithms for the parametric computation of mincost flows on graphs with n vertices where the output function has L breakpoints.

	Single commodity		
	Nondegenerate	Degenerate	Multi-commodity
Homogeneous	$\mathcal{O}(n^{2.375} + n^2 L)$	$\mathcal{O}(n^{2.375} + \mathrm{poly}(n)L)$	$\mathcal{O}(\mathrm{poly}(n)L)$
Nonhomogeneous	$\mathcal{O}(\mathrm{poly}(n) + n^2 L)$	$\mathcal{O}(\mathrm{poly}(n)L)$	

the same runtimes as for nonhomogenous functions even when the cost functions are homogeneous. Our results for multi-commodity networks can be generalized in a similar way.

Finally, we show that our algorithm is polynomial for single-commodity series-parallel networks. Specifically, we show that for these graphs, the number of breakpoints of the output function can be bounded by a polynomial in the input size of the graph. We also give an example of a single-commodity non–series-parallel network where the number of breakpoints of the output function is exponential in the input size of the graph.

1.2. Related Work

Potential flows as a model for gas, water, and electric network have a long history. The first mathematical treatment of these networks goes back to Birkhoff and Diaz [8], who show that (under reasonable assumptions) every demand vector imposes a unique flow. Similar results are obtained by Kirchhoff [34] and Duffin [16]. For the special case that the marginal cost functions f_e are linear (and not piecewise linear), the nonparametric version of the mincost flow problem (1) can be solved with quadratic programming techniques. Frank-Wolfe decomposition (Frank and Wolfe [18]) is based on solving linearized version of the objective and converges to the optimal solution. Lemke's algorithm (Sacher [52]) treats the problem as a linear complementarity problem. The Ellipsoid method (Kozlov et al. [38]) solves convex quadratic optimization problems in polynomial time. Kojima et al. [37] propose an extension of Lemke's algorithm for piecewise linear complementarity problems that is also able to solve the nonparametric variant of (1) for piecewise linear marginal cost functions f_e. It is worth noting that none of these approaches are able to solve the parametric problem studied in this paper. Specifically, Lemke's method introduces an artificial variable that is then traced continuously until an optimal solution is found. Intermediate solutions of this algorithm, however, are not optimal solutions of (1) for a smaller demand. Katzenelson [31] proposes an algorithm where demands are increased in a similar way as in our paper. However, when encountering a degenerate point, the algorithm jumps to another flow so that his algorithm does not solve the parametric problem.

The equilibrium principle for traffic flow is from Wardrop [57]. Beckmann et al. [6] characterize Wardrop equilibria in terms of the maximum of a convex program. Wardrop equilibria can be computed by the Frank-Wolfe method (Frank and Wolfe [18]). For refinements and variants of the Frank-Wolfe method for traffic equilibria, see Arezki and van Vliet [4], Bar-Gera [5], LeBlanc et al. [40], and Correa and Stier-Moses [14] for a survey. The computation of Wardrop equilibria as a function of the demand is an important intermediate step when computing optimal tolls for a subset of edges (Harks et al. [26]). They consider networks of parallel edges only and even for this special case, the flow on an edge as a function of the demand is governed by a partial differential equation where the second derivatives $f_e'(x_e)$ appear in the denominator (Harks et al. [26, lemma 4.3]). Restricting ourselves to piecewise quadratic functions avoids dealing with such rather intricate differential equations and allows to focus on the combinatorics of the problem, that is, the relevant supports.

In mathematical terms, our algorithms are homotopy methods. Homotopy methods are important in game theory, most notably the Lemke-Howson algorithm for computing equilibria in two-player games (Lemke and Howson [41]). For n-player games, the underlying problem becomes nonlinear. Herings and von den Elzen [29] propose an algorithm that approximates an equilibrium by solving piecewise linear approximations of the underlying complementarity problems (see Herings and Peeters [28] for a survey of further homotopy methods in game theory and Goldberg et al. [21] for a discussion of their complexity).

The standard mincost flow problem on directed graphs with linear costs and edge capacities $u_e \in \mathbb{R}_{\geq 0}$ is a special case of our model where the marginal cost functions are of the form $f_e(x) = -\infty$ for $x < 0$, $f_e(x) = +\infty$ for $x > u_e$, and $f_e(x) = a_e$ for some $a_e \in \mathbb{R}_{\geq 0}$. The standard successive shortest path algorithm implicitly solves the parametric variant of the problem. Zadeh [59] shows that the output of this parametric computation may be exponential in the input. Disser an Skutella [15] show that the successive shortest path algorithm is NP-mighty; that is, it can be used to solve NP-hard problems while being executed. When applied to these particular instances, our algorithm is similar to the successive shortest paths algorithm with the only difference being that all successive shortest paths are augmented at the same time. Because the result of Disser and Skutella on the NP-mightiness holds even when restricted to instances where there is always a unique shortest path, their result implies that our algorithm is NP-mighty as well. Ahuja et al. [1] propose an algorithm for the parametric variant of (1) where the marginal cost functions are piecewise constant and show that it is output-polynomial.

Further related is the parametric maximum flow problem with parametric edge capacities introduced by Gallo et al. [19]. Here, there is no cost function, and the capacity of some arcs depends on a single parameter; hence, the problem is unrelated to our parametric flow problems. Parametric variants have also been studied for other combinatorial problems such as cuts (Aissi et al. [3], Granot et al. [22]), shortest paths (Carstensen [10], Karp and Orlin [30], Nikolova et al. [48]), and linear programs (Murty [46]). Finally, we note that our algorithms can also be seen in the context of general parametric quadratic programs. For these problems, there exists an active set

method from Best [7], but the theoretical analysis assumes that the objective is quadratic (and not piecewise quadratic as in our setting), the constraint matrix has full rank (where the Laplace matrix in our setting has not full rank), and that the problem is nondegenerate (which we do not assume). Even under these assumptions, Best shows only that the active set methods terminates, whereas we can give concrete bounds on the runtime.

2. Preliminaries

We consider directed graphs $G = (V, E)$ with vertex set $V := \{1, \dots, n\}$ and edge set $E := \{e_1, \dots, e_m\} \subseteq V \times V$. We do not allow for multiple edges between two vertices, but note that these can be handled by introducing a dummy vertex in the middle of each of the multiple edges. We assume that G is strongly connected and encode G by its incidence matrix $\mathbf{\Gamma} = (\gamma_{v,e}) \in \mathbb{R}^{n \times m}$ defined as $\gamma_{v,e} = 1$, if edge e enters vertex v, $\gamma_{v,e} = -1$, if edge e leaves vertex v, and $\gamma_{v,e} = 0$, otherwise. For an edge $e = e_i$ for some $i \in \{1, \dots, m\}$, we denote by $\boldsymbol{\gamma}_e$ the ith column of $\mathbf{\Gamma}$. Throughout this work, all vectors are column vectors and will be denoted in bold.

Every vertex $v \in V$ has a demand $q_v \in R$ where a positive value $q_v > 0$ means that there is a demand of q_v flow units at vertex v, and a negative value $q_v < 0$ means that there is a supply of $|q_v|$ flow units at vertex v. We denote by $\mathbf{q} = (q_v)_{v \in V}$ the vector of all demands and assume that $\sum_{v \in V} q_v = 0$. Furthermore, every edge $e \in E$ is equipped with a strictly convex, continuous, piecewise quadratic cost function $F_e : \mathbb{R} \to \mathbb{R}$. More specifically, we assume that for every $e \in E$ there is a strictly increasing and piecewise linear *marginal cost function* $f_e : \mathbb{R} \to \mathbb{R}$ such that $F_e(x) = \int_0^x f_e(s)\,\mathrm{d}s$ for all $x \in \mathbb{R}$. For a function $f : \mathbb{R} \to \mathbb{R}$ and $x \in \mathbb{R}$, we denote the left-hand limit of f at x by $f^-(x) := \lim_{\xi \uparrow x} f(\xi)$, and the right-hand limit by $f^+(x) := \lim_{\xi \downarrow x} f(\xi)$.

A minimum cost flow for demand vector $\mathbf{q}$ is a vector $\mathbf{x} = (x_e)_{e \in E}$ solving

$$\min \sum_{e \in E} F_e(x_e) \quad \text{s.t.} \quad \mathbf{\Gamma x} = \mathbf{q}, \quad \mathbf{x} \geq 0. \tag{QF}$$

The following lemma expresses the optimality conditions of a solution to (QF).

Lemma 1. *A vector $\mathbf{x} \in \mathbb{R}^m$ is a minimum cost flow solving (QF) if and only if there is a vector $\boldsymbol{\pi} \in \mathbb{R}^n$ of vertex potentials with*

$$f_e^-(x_e) \leq \pi_w - \pi_v \leq f_e^+(x_e) \quad \text{for all } e = (v, w) \in E \text{ with } x_e > 0, \tag{3a}$$
$$\pi_w - \pi_v \leq f_e^+(x_e) \quad \text{for all } e = (v, w) \in E \text{ with } x_e = 0, \tag{3b}$$
$$\mathbf{\Gamma x} = \mathbf{q}, \tag{3c}$$
$$\mathbf{x} \geq 0. \tag{3d}$$

Proof. The minimization problem (QF) has affine-linear constraints only. The Karush-Kuhn-Tucker (KKT) conditions (Ruszczyński [51, theorem 3.34]) imply that $\mathbf{x}$ is an optimal solution of (QF) if and only if it is feasible and there exist dual variables $\boldsymbol{\pi} \in \mathbb{R}^n$ and $\boldsymbol{\mu} \in \mathbb{R}^m_{\geq 0}$ with $\boldsymbol{\mu}^\top \mathbf{x} = 0$ such that

$$0 \in \partial F(\mathbf{x}) + \mathbf{\Gamma}^\top \boldsymbol{\pi} - \boldsymbol{\mu}. \tag{4}$$

For every edge $e = (v, w) \in E$, this implies $0 = \partial F_e(x_e) + \pi_v - \pi_w - \mu_e$ where $\partial F_e(x_e) = [f_e^-(x_e), f_e^+(x_e)]$ is the subdifferential of F_e at x_e. This is equivalent to

$$\pi_w - \pi_v \in [f_e^-(x_e), f_e^+(x_e)] \qquad \text{if } x_e > 0, \tag{5a}$$
$$\pi_w - \pi_v \in [f_e^-(x_e) - \mu_e, f_e^+(x_e) - \mu_e] \qquad \text{if } x_e = 0. \tag{5b}$$

Let $\mathbf{x}$ be an optimal solution to (QF). Because $\mathbf{x}$ is feasible, it satisfies (3c) and (3d). Equation (5a) directly implies (3a). From (5b), we further obtain $\pi_w - \pi_v \leq f_e^+(x_e) - \mu_e \leq f_e^+(x_e)$ so that (3b) is satisfied.

Conversely, let $\mathbf{x}$ and $\boldsymbol{\pi}$ satisfy (3a)–(3d). Set $\mu_e = f_e^+(x_e) - (\pi_w - \pi_v)$ for all $e = (v, w) \in E$ with $x_e = 0$. then, because $f_e^-(x_e) - f_e^+(x_e) \leq 0$ as f_e is strictly increasing, we obtain (5b). □

2.1. Problem Statement and Variants

The main problem we address in this paper is the parametric computation of minimum cost flows for networks with piecewise quadratic cost. Formally, we consider the following problem:

PARAMETRIC MINCOST FLOW

GIVEN: graph $G = (V, E)$, demand vector $\mathbf{q}$, strictly convex, continuous, piecewise quadratic edge cost functions $F_e(x), e \in E$;
FIND: map $\mathbf{x} : \mathbb{R}_{\geq 0} \to \mathbb{R}^m$
s.t. $\mathbf{x}(\lambda) = (x_e(\lambda))_{e \in E}$ is a minimum cost flow for demands $\lambda \mathbf{q}$ for all $\lambda \geq 0$.

With the techniques developed in this paper, we can also solve a slightly more general parametric problem, where we are given a piecewise linear function $\mathbf{h}: \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ and we compute a map $\mathbf{x}: \mathbb{R}_{\geq 0} \to \mathbb{R}^m$ with the property that $\mathbf{x}(\lambda)$ is a minimum cost flow for demands $\mathbf{h}(\lambda)$ for all $\lambda \geq 0$. For ease of exposition, we discuss in this paper only the linear case where $\mathbf{h}(\lambda) = \lambda \mathbf{q}$ for some given $\mathbf{q} \in \mathbb{R}^n$.

Our techniques further allow to handle the case of upper and lower bounds on the flow on the edges. Formally, for each edge $e \in E$ let $l_e \in \mathbb{R}_{\geq 0}$ be a lower capacity and $u_e \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ be an upper capacity on the flow of edge e. In this more general model, a flow $\mathbf{x}$ is feasible if in addition to the constraints of (QF), the constraint $l_e \leq x_e \leq u_e$ is also satisfied for all $e \in E$. Not all demands $\lambda \mathbf{q}$ with $\lambda \geq 0$ may be feasible. However, it is easy to see that the maximal set

$$\Lambda = \{\lambda \in \mathbb{R}_{\geq 0} \mid \text{there is a feasible flow } \mathbf{x} \text{ satisfying } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\}$$

is convex and closed. The following more general problem asks for the computation of Λ and the mincost flows parametrized by $\lambda \in \Lambda$.

PARAMETRIC MINCOST FLOW WITH LOWER AND UPPER CAPACITIES

GIVEN: graph $G = (V, E)$, demand vector $\mathbf{q}$, lower capacities $\mathbf{l} \in \mathbb{R}^m_{\geq 0}$, upper capacities $\mathbf{u} \in (\mathbb{R}_{\geq 0} \cup \{\infty\})^m$
strictly convex, continuous, piecewise quadratic edge cost functions $F_e(x), e \in E$;
FIND: maximal set $\Lambda \subseteq \mathbb{R}_{\geq 0}$ s.t. there is a feasible flow for all demands in $\{\lambda \mathbf{q} \mid \lambda \in \Lambda\}$
map $\mathbf{x}: \Lambda \to \mathbb{R}^m$
s.t. $\mathbf{x}(\lambda) = (x_e(\lambda))_{e \in E}$ is a minimum cost flow for demands $\lambda \mathbf{q}$ for all $\lambda \in \Lambda$.

It is worth mentioning that the standard method to treat upper and lower bounds on the flow of each edge by modifying the graph (Ahuja et al. [2, section 2.4]) does not work in our setting because it introduces an edge with zero cost, whereas we assume that all cost functions are strictly convex.

We will show that the solution $\mathbf{x}(\lambda)$ to PARAMETRIC MINCOST FLOW is a piecewise linear function with finitely many breakpoints $\lambda_0, \lambda_1, \ldots, \lambda_{L-1}, \lambda_L$ (including the two artificial breakpoints $\lambda_0 = 0$ and $\lambda_L = \infty$) for some $L \in \mathbb{N}$; that is, for all $i \in \{1, \ldots, L-1\}$ and all $e \in E$ there are $\alpha_{e,i}, \beta_{e,i} \in \mathbb{R}$ such that $\mathbf{x}(\lambda) = (x_e(\lambda))_{e \in E}$ with

$$x_e(\lambda) = \alpha_{e,i} \lambda + \beta_{e,i} \quad \text{for all} \quad \lambda \in [\lambda_i, \lambda_{i+1}].$$

Therefore, to solve PARAMETRIC MINCOST FLOW, it is sufficient to compute the coefficients $\alpha_{e,i}$ and $\beta_{e,i}$ for all $e \in E$ and $i \in \{0, \ldots, L\}$. We show in Section 6 that the number of breakpoints L is always finite (as long as the number of breakpoints of the piecewise quadratic cost functions F_e is finite). For series-parallel graphs, L is also polynomially bounded in the size of the input, that is, the size of the graph and the number of breakpoints of the input functions. However, for general, non–series-parallel graphs, the number of breakpoints can be exponential in the size of the input. We discuss this in more detail in Section 6.

We illustrate the parametric computation of mincost flows with the following example.

Example 1. Consider the graph from Figure 1(a) with the given edge cost functions $F_e(x)$ on the edges. Assume that we are given demands $\mathbf{q} = (q_v)_{v \in V}$ with $q_s = -1$, $q_{v_1} = q_{v_2} = 0$, and $q_t = 1$. For $\lambda = 1$, any s-t flow of rate 1 is feasible. The mincost flow for this demand is $\mathbf{x} = (x_e)_{e \in E}$ with $x_e = 1$ for $e \in \{(s, v_1), (v_1, v_2), (v_2, t)\}$ and $x_e = 0$ otherwise. It may be easily verified that this is a mincost flow using the vertex potential $\boldsymbol{\pi} = (\pi_s, \pi_{v_1}, \pi_{v_2}, \pi_t)^\top = (0,2,3,5)^\top$ and the optimality conditions from Lemma 1.

Let us consider the PARAMETRIC MINCOST FLOW problem for this setting. In contrast to the fixed demand setting, we now aim to compute mincost flow for *all* demands $\lambda \mathbf{q}$; that is, in this example, mincost s-t flows with variable flow rate λ. The solution to this problem is a family of functions $\mathbf{x}(\lambda) = (x_e(\lambda))_{e \in E}$ such that $\mathbf{x}(\lambda)$ is an s-t flow of rate $\lambda \geq 0$. Figure 1(b) depicts the functions solving PARAMETRIC MINCOST FLOW for this example. We present the explicit computation of the solution to this problem in Example 5 at the end of Section 4. For every $\lambda \geq 0$, we can use the functions $\mathbf{x}(\lambda)$ to obtain the mincost flows the demand $\lambda \mathbf{q}$. In Figure 1(c), we see three examples of mincost flow for different λ that can be determined from the flow functions in Figure 1(b). The numbers next to the vertices are the corresponding optimal potentials.

2.2. Outline and Organization of the Paper

Our main tool to solve the parametric mincost flow problem is a relaxation of the problem where the nonnegativity constraint in (QF) is omitted. The relaxation can be interpreted as an electrical flow by associating the flow with the electrical current and the vertex potentials with the voltage. The slope of the piecewise linear cost

Figure 1. Input and solution of the parametric mincost flow problem discussed in Example 1.

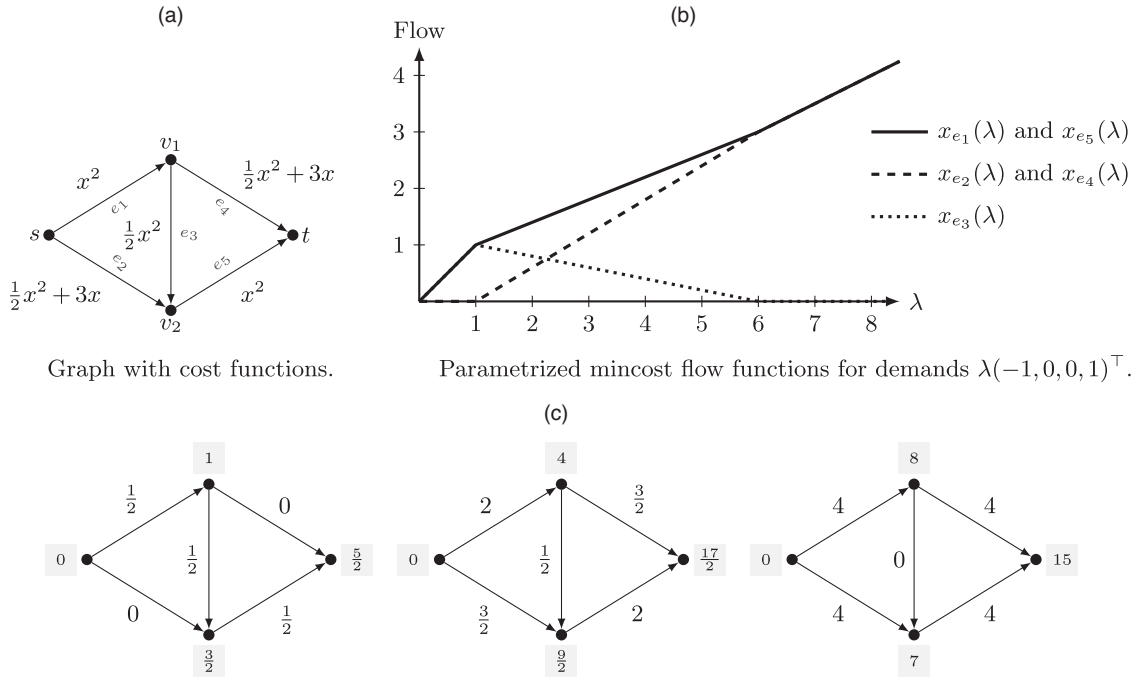

Graph with cost functions. Parametrized mincost flow functions for demands $\lambda(-1,0,0,1)^\top$.

Mincost flows for $\lambda = \frac{1}{2}$, $\lambda = \frac{7}{2}$, $\lambda = 8$. Flow on edges, optimal potentials next to the vertices.

functions then corresponds to the resistance of the edge. In light of this connection, we call the relaxation an *electrical flow*.

In Section 3, we study the parametric computation of electrical flows for the special case where all marginal cost functions f_e are continuous and homogeneous (i.e., $f_e(0) = 0$). For an electrical network, the continuity of the marginal cost function implies that the electrical current along an edge depends continuously on the differences of the voltages of its end vertices. This is a reasonable assumption for electrical networks as long as there are no special components such as transistors or diodes. The homogeneity of the marginal cost functions implies that the all-zero flow can be realized by an all-zero voltage. This is sensible as long as there are no active components on the edges that produce power. Under these two assumptions, we propose an algorithm that computes parametric electrical flows $\mathbf{x}(\lambda)$ for all $\lambda \in \mathbb{R}_{\geq 0}$ such that $\mathbf{x}(\lambda)$ satisfies the demands $\lambda\mathbf{q}$ for all λ. Our algorithm starts with the all-zero potential vector that induces the all-zero electrical flow by assumption. Then, the parameter λ is increased continuously. As long as the support of the flow does not change, there is a linear and bijective relationship between vertex potentials and flows that can be expressed by a weighted graph Laplacian (where the weights correspond to the conductivities of the support). When a new edge enters or leaves the support, a new Laplacian can be computed by a simple and computational inexpensive rank 1 update. The case where multiple edges enter or leave the support at the same λ can be treated by an implicit perturbation technique similar to perturbation techniques for the Simplex algorithm. This yields an algorithm that is output-polynomial when no perturbation is needed.

In Section 4, we generalize this method to the case of discontinuous and nonhomogeneous marginal cost functions f_e. For electrical networks, discontinuous marginal cost functions allow to model special components such as diodes or transistors that let only special flows pass. More importantly, this extended method is able to solve Parametric Mincost Flow by replacing all original marginal cost functions f_e by new marginal cost functions $\tilde{f}_e$ defined by $\tilde{f}_e(x) = f_e(x)$ for $x \geq 0$ and $\tilde{f}_e(x) = -\infty$ for $x < 0$. The new marginal cost functions $\tilde{f}_e$ are discontinuous and essentially act as diodes that let only flow in one direction pass. A similar construction also allows to model upper bounds u_e on the flow on edge e. To treat also lower bounds l_e on the flow on an edge, we allow also for nonhomogeneous cost functions f_e. Mathematically, this is challenging because the all-zero potential vector may then not be feasible for the all-zero flow. We circumvent this issue by extending our algorithm by an additional phase that computes a feasible potential vector for the all-zero flow. Being able to treat noncontinuous marginal cost functions not necessarily satisfying $f_e(0) = 0$ thus allows us to solve Parametric Mincost Flow with Lower and Upper Capacities.

In Section 5, we generalize our method to multi-commodity mincost flows.

Finally, in Section 6, we show that our method runs, except for degenerate cases, in output-polynomial time; that is, it runs in polynomial time in the size of the input and the output function $\mathbf{x}$, where the latter is determined by the number of its breakpoints L.

3. Electrical Flows with Homogeneous and Continuous Marginal Costs

Relaxing the non-negativity constraint of the piecewise quadratic mincost flow problem (QF), we obtain the problem

$$\min \quad \sum_{e\in E} F_e(x_e) \quad \text{s.t.} \quad \mathbf{\Gamma x} = \mathbf{q}. \tag{EF}$$

This problem is still defined on a directed graph G, but for all edges both positive flow values (going in the direction of the edge) and negative flow values (going in the opposite direction) is allowed.

This setting can be interpreted as an electrical flow in the following way. In the linear case where $f_e(x) = ax$ with $a > 0$, the marginal cost function represents Ohm's law stating that the voltage along an electrical conductor is proportional to the electrical current on this edge, where the slope a is the constant resistance of the conductor. In this section, we assume that the marginal cost function f_e is continuous, strictly increasing, piecewise linear, and homogeneous (i.e., $f_e(0) = 0$). The physical interpretation of the latter condition is that we only allow passive components (i.e., power consuming components such as resistors or motors) and exclude active components (i.e., power generating components such as batteries). We will generalize our setting to discontinuous and non-homogeneous marginal cost functions in Section 4.

The following characterization of electrical flows for continuous cost functions is a direct corollary of Lemma 1 after relaxing the non-negativity constraints.

Corollary 1. *A vector $\mathbf{x} \in \mathbb{R}^m$ is an electrical flow solving (EF) for continuous marginal cost functions f_e, $e \in E$ if and only if there is a vector $\boldsymbol{\pi} \in \mathbb{R}^n$ of vertex potentials such that*

$$\pi_w - \pi_v = f_e(x_e) \quad \text{for all } e = (v,w) \in E, \tag{6a}$$

$$\mathbf{\Gamma x} = \mathbf{q}. \tag{6b}$$

Proof. The proof follows from the proof of Lemma 1 using that the non-negativity conditions, the dual multipliers μ_e are relaxed and that the marginal cost functions are continuous, and hence, $f_e^+(x) = f_e^-(x)$ for all $x \in \mathbb{R}$, $e \in E$. □

The optimality conditions from Lemma 1 are known as Ohm's Law (6a) and Kirchhoff's law (6b). In this section, we are interested in the following relaxation of Parametric Mincost Flow.

Parametric Electrical Flow

Given: graph $G = (V,E)$, demand vector $\mathbf{q}$, strictly increasing, piecewise linear marginal cost functions $f_e(x), e \in E$;
Find: map $\mathbf{x} : \mathbb{R} \to \mathbb{R}^m$
s.t. $\mathbf{x}(\lambda) = (x_e(\lambda))_{e\in E}$ is an electrical flow for demands $\lambda\mathbf{q}$ for all $\lambda \geq 0$.

3.1. Structure of Electrical Flows

In order to compute an electrical flow, it is sufficient to solve the system (6a)-(6b). Note that (6a) can be solved uniquely for x_e as we assume that the functions f_e are strictly increasing. The flow $\mathbf{x}$ is determined by the differences in the potential $\pi_w - \pi_v$ along edges rather than the absolute values. In order to obtain unique solutions, we fix the potential at some vertex $s \in V$. We use the convention that the potential of the vertex $s := v_1$ is normalized to zero, that is, $\pi_s := 0$. From now on, we only consider potentials in the subspace

$$\Pi := \{\boldsymbol{\pi} \in \mathbb{R}^n \mid \pi_s = 0\} \tag{7}$$

of $\mathbb{R}^n$ containing all potential vectors that are normalized to zero at the vertex $s \in V$. We refer to Π as the *potential space*.

In the case of linear marginal cost functions f_e, the system (6a)-(6b) is a system of linear equations and can be solved directly. Because we assume that the function f_e is piecewise linear, it will be convenient to consider

regions in the space of all potentials Π where the functions f_e are linear. More precisely, we define regions $R \subseteq \Pi$ such that the function

$$\mathbf{f}^{-1} : \Pi \to \mathbb{R}^m \qquad \boldsymbol{\pi} \mapsto \mathbf{f}^{-1}(\boldsymbol{\pi}) := (f_e^{-1}(\pi_w - \pi_v))_{e \in E}$$

is a linear function for all $\boldsymbol{\pi} \in R$. Given the correct region, the system (6a)-(6b) is again a linear system that can be solved easily.

To define the regions, recall that for every edge e there is a set $T_e = \{1, \dots, \bar{t}_e\}$ and $\bar{t}_e + 1$ breakpoints $\tau_{e,t} \in \mathbb{R} \cup \{-\infty, +\infty\}$, $t \in \{1, \dots, \bar{t}_e + 1\}$ such that $-\infty = \tau_{e,1} < \tau_{e,2} < \dots < \tau_{e,\bar{t}_e} < \tau_{e,\bar{t}_e+1} = \infty$. Furthermore, for all $t \in T_e$, there are constants $a_{e,t} > 0$ and $b_{e,t} \in \mathbb{R}$ such that $f_e(x) = a_{e,t}\, x + b_{e,t}$ for all $x \in [\tau_{e,t}, \tau_{e,t+1}]$. Let $\sigma_{e,t} := f_e(\tau_{e,t})$ be the value of f_e at every breakpoint $\tau_{e,t}$ with the convention that $f_e(-\infty) := -\infty$ and $f(+\infty) := +\infty$. Furthermore, let $\bar{T} := T_{e_1} \times \dots \times T_{e_m}$.

Then, for every $\mathbf{t} \in \bar{T}$, we refer to the polytope

$$R_{\mathbf{t}} := \{\boldsymbol{\pi} \in \Pi \mid \sigma_{e,t_e} \le \boldsymbol{\gamma}_e^\top \boldsymbol{\pi} \le \sigma_{e,t_e+1} \text{ for all edges } e \in E\} \tag{8}$$

as the *region in the potential space* corresponding to $\mathbf{t} \in \bar{T}$. Considering regions in the potential space has the advantage that for every $\boldsymbol{\pi} \in R_{\mathbf{t}}$. Equation (6a) has the explicit solution

$$x_e = f_e^{-1}(\pi_w - \pi_v) = \frac{1}{a_{e,t_e}}(\pi_w - \pi_v) - \frac{b_{e,t_e}}{a_{e,t_e}}$$

for all $e \in E$. In particular, we obtain a closed form for the function $\mathbf{f}^{-1}$ for all potentials $\boldsymbol{\pi} \in R_{\mathbf{t}}$. For an illustration of these concepts, we give the following example.

Example 2. Consider the graph from Figure 2(a) with piecewise linear marginal cost functions on the edges. Every marginal cost function has three breakpoints (including the artificial breakpoints $-\infty, \infty$), that is, $T_e = \{1,2\}$. The set $\bar{T} = \{1,2\}^3$ contains eight possible combinations. Thus, there are eight regions $R_{\mathbf{t}}, \mathbf{t} \in \bar{T}$. As an example, consider the region $R_{\mathbf{t}}$ with $\mathbf{t} = (1,1,2)^\top$. This is the set of all potentials belonging to flows with $x_{e_1} \le 3$ and $x_{e_2} \le 2$ and $x_{e_3} \ge 1$. The respective inequalities in (8) read $\pi_v - \pi_s \le \sigma_{e_1,2} = f_{e_1}(3) = 3$ for the edge e_1, $\pi_t - \pi_v \le \sigma_{e_2,2} = f_{e_2}(2) = 2$ for the edge e_2, and $\pi_t - \pi_s \ge \sigma_{e_2,2} = f_{e_3}(1) = 1$. Figure 2(b) shows all region polytopes in the potential space for fixed $\pi_s = 0$. There are only seven regions in Figure 2(b) because the region $R_{(2,2,1)^\top}$ is empty as the respective inequalities in (8) are inconsistent. By definition, the function $\mathbf{f}^{-1}$ has a closed, linear form in every region. In particular, for region $R_{(1,1,2)^\top}$, the function has the explicit linear form $\mathbf{f}^{-1}(\boldsymbol{\pi}) = (\pi_v - \pi_s, \pi_t - \pi_v, \frac{1}{4}(\pi_t - \pi_s + 3))$ for all potentials $\boldsymbol{\pi} = (\pi_s, \pi_v, \pi_t)^\top \in R_{(1,1,2)^\top}$. Hence, for example, the potential $\boldsymbol{\pi} = (0,1,2)^\top \in R_{(1,1,2)^\top}$ induces the flow $\mathbf{x} = \mathbf{f}^{-1}(\boldsymbol{\pi}) = (1,1,\frac{5}{4})^\top$, meaning that this flow $\mathbf{x}$ together with the potential $\boldsymbol{\pi}$ satisfy the optimality conditions of Lemma 1.

Figure 2. (a) A graph with piecewise linear marginal cost functions. (b) The potential space Π for fixed $\pi_s = 0$. The breakpoints of the cost functions induce boundary hyperplanes (dashed lines, with normal vectors $\boldsymbol{\gamma}_e$) that subdivide the space into regions.

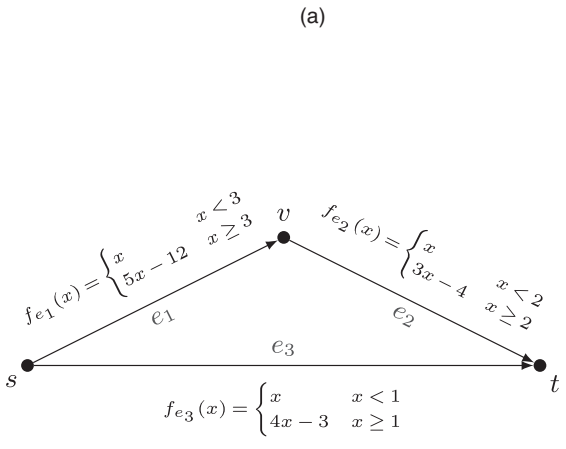

Graph with marginal cost functions.

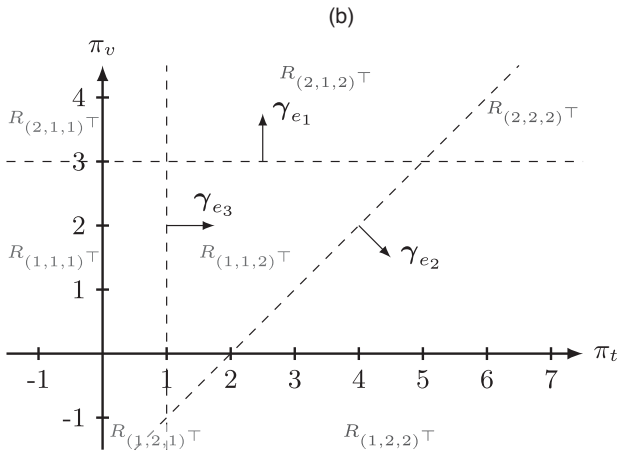

Regions in the potential space Π.

The following lemma gives an explicit matrix form for the induced flow function $\mathbf{f}^{-1}$ inside a region $R_{\mathbf{t}}$.

Lemma 2. *The function $\mathbf{f}^{-1}: \Pi \to \mathbb{R}^m$ is a well-defined, piecewise linear, and injective function. For any $\mathbf{t} \in \bar{T}$, the induced flow $\mathbf{f}^{-1}(\boldsymbol{\pi})$ of any potential $\boldsymbol{\pi} \in R_{\mathbf{t}}$ can be expressed as*

$$\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^{\top}\boldsymbol{\pi} - \mathbf{d}_{\mathbf{t}},$$

where $\mathbf{C}_{\mathbf{t}}$ is the diagonal matrix $\mathbf{C}_{\mathbf{t}} = \operatorname{diag}(c_{e_1,t_{e_1}}, c_{e_2,t_{e_2}}, \ldots, c_{e_m,t_{e_m}}) = \operatorname{diag}(1/a_{e_1,t_{e_1}}, \ldots, 1/a_{e_m,t_{e_m}})$ of conductances and $\mathbf{d}_{\mathbf{t}}$ is the vector $\mathbf{d}_{\mathbf{t}} = (d_{e_1,t_{e_1}}, \ldots, d_{e_m,t_{e_m}})^{\top} = (b_{1,t_1}/a_{1,t_1}, \ldots, b_{e_m,t_{e_m}}/a_{e_m,t_{e_m}})^{\top}$.

Proof. Given a potential $\boldsymbol{\pi} \in \Pi$, its induced flow satisfies $\pi_w - \pi_v = f_e(x_e)$. Because the marginal cost functions f_e are strictly increasing and piecewise linear, their inverse f_e^{-1} exists and is piecewise linear as well. Hence, the first part follows. The second part follows immediately from the definition of the region $R_{\mathbf{t}}$ and the closed form solution of (6a) for all potentials $\boldsymbol{\pi} \in R_{\mathbf{t}}$. □

Lemma 2 shows that there is a one-to-one correspondence between potentials in Π and potential based flows (i.e., flows satisfying (6a)). Hence, it will be sufficient to compute the potentials $\boldsymbol{\pi}$ of a solution—the explicit function $\mathbf{f}^{-1}$ yields the corresponding electrical flow. A potential $\boldsymbol{\pi}$ is a potential of an electrical flow, if its induced flow $\mathbf{x} = \mathbf{f}^{-1}(\boldsymbol{\pi})$ satisfies (6b) as well. Thus, with the explicit form from Lemma 2, we obtain that given demands $\mathbf{q}$, a potential $\boldsymbol{\pi} \in R_{\mathbf{t}}$ is a potential of an electrical flow (and, hence, the flow $\mathbf{x} = \mathbf{f}^{-1}(\boldsymbol{\pi})$ is an electrical flow) if and only if

$$\mathbf{L}_{\mathbf{t}}\boldsymbol{\pi} - \tilde{\mathbf{d}}_{\mathbf{t}} = \mathbf{q}, \tag{9}$$

where $\mathbf{L}_{\mathbf{t}} := \boldsymbol{\Gamma}\mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^{\top}$ and $\tilde{\mathbf{d}}_{\mathbf{t}} := \boldsymbol{\Gamma}\mathbf{d}_{\mathbf{t}}$.

Definition 1. The matrix $\mathbf{L}_{\mathbf{t}} := \boldsymbol{\Gamma}\mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^{\top}$ is called the weighted Laplacian matrix of G in $R_{\mathbf{t}}$.

The following lemma states some well-known properties of Laplacians of weighted graphs (Grone [23], Merris [42], Mohar et al. [45]).

Lemma 3. *For all $\mathbf{t} \in \bar{T}$, the Laplacian $\mathbf{L}_{\mathbf{t}} = \boldsymbol{\Gamma}\mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^{\top}$ has the following properties:*

i. *The rank of $\mathbf{L}_{\mathbf{t}}$ is $n - n_C$ where n_C is the number of connected components and two vertices are considered connected if they are connected by an edge with $c_{e,t_e} \neq 0$.*

ii. *The matrix $\mathbf{L}_{\mathbf{t}}$ is positive semi-definite.*

iii. *The row sum and column sum of $\mathbf{L}_{\mathbf{t}}$ is zero for every row or column.*

iv. *The term f G is connected then the reduced Laplacian matrix $\hat{\mathbf{L}}_{\mathbf{t}}$ obtained from the matrix $\mathbf{L}_{\mathbf{t}}$ by deleting the first row and column is positive definite and inverse-positive, i.e., $\mathbf{L}_{\mathbf{t}}^{-1} \geq 0$.*

From Lemma 3(i) and (iii), we infer that the nullspace of the Laplacian matrix $\mathbf{L}_{\mathbf{t}}$ is spanned by the all-ones vector $\mathbf{1} \in \mathbb{R}^n$, that is, $\ker(\mathbf{L}_{\mathbf{t}}) = \operatorname{span}(\mathbf{1})$. Thus, the mapping $\boldsymbol{\pi} \longmapsto \mathbf{L}_{\mathbf{t}}\boldsymbol{\pi} - \tilde{\mathbf{d}}_{\mathbf{t}}$ is injective on the subspace Π, and hence, (9) has a unique solution $\boldsymbol{\pi}$ in Π for a fixed demand vector $\mathbf{q}$. In particular, there is a unique generalized inverse of $\mathbf{L}_{\mathbf{t}}$ (that is a matrix $\mathbf{L}_{\mathbf{t}}^*$ with $\mathbf{L}_{\mathbf{t}}\mathbf{L}_{\mathbf{t}}^*\mathbf{L}_{\mathbf{t}} = \mathbf{L}_{\mathbf{t}}$) that maps into Π. This unique generalized inverse is

$$\mathbf{L}_{\mathbf{t}}^* := \begin{pmatrix} 0 & \mathbf{0}^{\top} \\ \mathbf{0} & \hat{\mathbf{L}}^{-1} \end{pmatrix},$$

where $\hat{\mathbf{L}}$ is the matrix obtained from $\mathbf{L}_{\mathbf{t}}$ by deleting the first row and column. This matrix maps into Π by definition and it can be easily checked that it is a generalized inverse of $\mathbf{L}_{\mathbf{t}}$. Overall, *given the correct region $R_{\mathbf{t}}$*, the electrical flow solving the system (6a)-(6b) is

$$\mathbf{x} = \mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\mathbf{L}_{\mathbf{t}}^*(\mathbf{q} + \tilde{\mathbf{d}}_{\mathbf{t}})) = \mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}\mathbf{L}_{\mathbf{t}}^*(\mathbf{q} + \tilde{\mathbf{d}}_{\mathbf{t}}) - \mathbf{d}_{\mathbf{t}}.$$

The central part of computing an electrical flow in the case of piecewise linear marginal cost functions f_e therefore lies in finding the correct region $R_{\mathbf{t}}$ such that the potential of the solution lies in that region.

3.2. Structure of Parametric Electrical Flows

Because the objective function in (EF) is strictly convex, there is a unique eletrical flow for every demand vector $\mathbf{q}$ with $\sum_{v \in V} q_v = 0$. The function $\mathbf{f}: \mathbb{R}^m \to \Pi$ is injective; thus, there is also a unique potential $\boldsymbol{\pi} \in \Pi$ for every demand vector $\mathbf{q}$. Hence, given some demand vector $\mathbf{q}$, the function $\lambda \longmapsto \boldsymbol{\pi}(\lambda)$ with $\boldsymbol{\pi}(\lambda)$ being the unique potential in Π of the electrical flow for demand $\lambda\mathbf{q}$, $\lambda \geq 0$ is well defined. We refer to this function $\lambda \longmapsto \boldsymbol{\pi}(\lambda)$ as the *solution curve in the potential space* of the Parametric Electrical Flow problem. Given the solution curve $\boldsymbol{\pi}(\cdot)$, the function

$$\mathbf{x}: \mathbb{R}_{\geq 0} \to \mathbb{R}^m, \quad \lambda \longmapsto \mathbf{f}^{-1}(\boldsymbol{\pi}(\lambda))$$

Figure 3. The solution curve and the electrical flow function in the network from Example 2 for the demand $\mathbf{q} = (q_s, q_v, q_t)^\top = (-1, 0, 1)$, that is, for s-t flows. (a) The solution curve (thick line) inside the potential space and regions with potential directions $\Delta\pi_{\mathbf{t}}$. (b) The electrical flow functions for this example $\lambda \longmapsto \mathbf{x}(\lambda) = \mathbf{f}^{-1}(\boldsymbol{\pi}(\lambda))$.

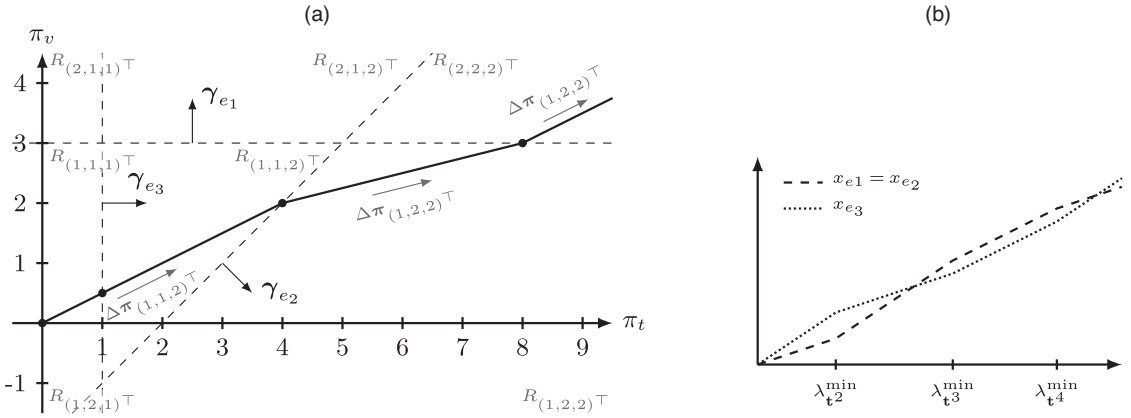

Regions and solution curve $\boldsymbol{\pi}(\cdot)$ in the potential space Π. Electrical flow functions $\mathbf{x}(\cdot)$.

is the (unique) solution to Parametric Electrical Flow. We have shown that the computation of electrical flows and their potentials (and thus the computation of the solution curve) depends on finding the correct regions in which the solution (curve) lies. To this end, we consider the sets

$$\Pi_{\mathbf{t}} = \{\boldsymbol{\pi} \in R_{\mathbf{t}} \mid \exists \lambda \geq 0 \text{ with } \mathbf{L}_{\mathbf{t}}\boldsymbol{\pi} - \tilde{\mathbf{d}}_{\mathbf{t}} = \lambda \mathbf{q}\}$$

and refer to $\Pi_{\mathbf{t}}$ as the *solution segment in region* $R_{\mathbf{t}}$. The set $\Pi_{\mathbf{t}}$ contains all potentials belonging to the image of the solution curve that lie in the region $R_{\mathbf{t}}$. Some regions may not contain any potential of the image of $\boldsymbol{\pi}(\cdot)$, then $\Pi_{\mathbf{t}} = \emptyset$ and we call these regions *infeasible regions*. For every region, we define the *potential direction* $\Delta\boldsymbol{\pi}_{\mathbf{t}} := \mathbf{L}^*_{\mathbf{t}}\mathbf{q}$ and the *potential offset* $\boldsymbol{\pi}_{\mathbf{t}} := \mathbf{L}^*_{\mathbf{t}}\tilde{\mathbf{d}}_{\mathbf{t}}$. Together with the definition of the region $R_{\mathbf{t}}$, we then obtain

$$\Pi_{\mathbf{t}} = \{\boldsymbol{\pi}_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}} \mid \exists \lambda \geq 0 \text{ with } \sigma_{e,t_e} \leq \boldsymbol{\gamma}_e^\top(\boldsymbol{\pi}_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}}) \leq \sigma_{e,t_e+1} \text{ for all edges } e \in E\}. \tag{10}$$

Thus, the solution segment in region $R_{\mathbf{t}}$ is a (possibly empty) segment of the line $\boldsymbol{\pi}_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}}$ (see Figure 3(a) for a graphical representation of the segments of the solution curve). The inequalities in (10) can be solved for λ, yielding an explicit lower and upper bound on λ for which $\boldsymbol{\pi}_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}}$ is part of the solution curve.

Theorem 1. *The image of the solution curve is the union of the solution segments, that is,*

$$\boldsymbol{\pi}(\mathbb{R}_{\geq 0}) := \{\boldsymbol{\pi}(\lambda) \mid \lambda \geq 0\} = \bigcup_{\mathbf{t} \in \bar{T}} \Pi_{\mathbf{t}}.$$

Furthermore, for every region $R_{\mathbf{t}}$, there are numbers $\lambda^{\min}_{\mathbf{t}}$ and $\lambda^{\max}_{\mathbf{t}}$ such that

$$\Pi_{\mathbf{t}} = \{\boldsymbol{\pi}_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}} \mid \lambda \in [\lambda^{\min}_{\mathbf{t}}, \lambda^{\max}_{\mathbf{t}}]\}. \tag{11}$$

Proof. The first part follows directly from the definition of the solution segments $\Pi_{\mathbf{t}}$. If $\boldsymbol{\pi} \in \Pi_{\mathbf{t}}$ for some $\mathbf{t}$ then by definition its induced flow is an electrical flow for demand $\lambda\mathbf{q}$. On the other hand, if $\boldsymbol{\pi} = \boldsymbol{\pi}(\lambda)$ for some $\lambda \geq 0$, then $\boldsymbol{\pi} \in R_{\mathbf{t}}$ for some $\mathbf{t}$. However, then it must be also in $\Pi_{\mathbf{t}}$ by definition.

For the second part, let $\mathbf{t} \in \bar{T}$ and define the values

$$\underline{\lambda}_{\mathbf{t}}(e) := \frac{\sigma_{e,t_e} - \boldsymbol{\gamma}_e^\top\boldsymbol{\pi}_{\mathbf{t}}}{\boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}}}$$

$$\overline{\lambda}_{\mathbf{t}}(e) := \frac{\sigma_{e,t_e+1} - \boldsymbol{\gamma}_e^\top\boldsymbol{\pi}_{\mathbf{t}}}{\boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}}}$$

for every edge $e \in E$ with $\boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}} \neq 0$. The values $\underline{\lambda}_{\mathbf{t}}(e)$ and $\overline{\lambda}_{\mathbf{t}}(e)$ are the λ values for which the inequalities in (10) are satisfied with equality. Defining

$$\lambda^{\min}_{\mathbf{t}} := \max\left(\{\underline{\lambda}_{\mathbf{t}}(e) \mid \boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}} > 0\} \cup \{\overline{\lambda}_{\mathbf{t}}(e) \mid \boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}} < 0\} \cup \{0\}\right)$$

$$\lambda^{\max}_{\mathbf{t}} := \min\left(\{\underline{\lambda}_{\mathbf{t}}(e) \mid \boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}} < 0\} \cup \{\overline{\lambda}_{\mathbf{t}}(e) \mid \boldsymbol{\gamma}_e^\top\Delta\boldsymbol{\pi}_{\mathbf{t}} > 0\}\right)$$

yields (11). □

We note that for infeasible regions, we have $\lambda_{\mathbf{t}}^{\min} > \lambda_{\mathbf{t}}^{\max}$ yielding an empty interval $[\lambda_{\mathbf{t}}^{\min}, \lambda_{\mathbf{t}}^{\max}]$. Furthermore, it is possible that $\lambda_{\mathbf{t}}^{\max} = \infty$ for some regions. In these regions, the solution curve can be extended to infinity.

The results of this section and Theorem 1 can be summarized as follows. In order to solve Parametric Electrical Flows, it is sufficient to compute the solution curve $\boldsymbol{\pi}(\cdot)$ as the equilibrium flows can be obtained from it using the induced flow function $\mathbf{f}^{-1}$. The solution curve itself consists of linear line segments in different regions of the potential space Π. Every of these line segments has the form stated in Theorem 1 with a fixed domain $[\lambda_{\mathbf{t}}^{\min}, \lambda_{\mathbf{t}}^{\max}]$. This domain as well as the offset $\pi_{\mathbf{t}}$ and the direction vector $\Delta\boldsymbol{\pi}_{\mathbf{t}}$ that define the line segment in every region can be computed with the generalized inverse of the Laplacian matrix $\mathbf{L}_{\mathbf{t}}^*$.

Figure 3 depicts the piecewise linear solution curve $\lambda \mapsto \boldsymbol{\pi}(\lambda)$ inside the potential space (Figure 3(a)) and the electrical flow functions (Figure 3(b)) solving Parametric Electrical Flows $\lambda \mapsto \mathbf{x}(\lambda) = \mathbf{f}^{-1}(\boldsymbol{\pi}(\lambda))$ for the graph from Example 2 for the demand $\mathbf{q} = (q_s, q_v, q_t)^\top = (-1, 0, 1)^\top$.

3.3. Basic Algorithm for Electrical Flows

We are now ready to describe our basic algorithm solving Parametric Electrical Flow. As shown in the last section, it suffices to compute the solution curve $\boldsymbol{\pi}(\cdot)$ in the potential space. The latter consists of piecewise linear line segments that we compute with our algorithm. The basic idea is the following: We start with the zero flow $\mathbf{x}(0) = \mathbf{0}$ for $\lambda = 0$. Because $f_e(0) = 0$, this flow is induced by the potential $\boldsymbol{\pi}^0 = \mathbf{0}$. Then, we determine $\mathbf{t}$ such that $\boldsymbol{\pi}^0 \in R_{\mathbf{t}}$. By Theorem 1, this allows us to compute the whole line segment $\Pi_{\mathbf{t}}$ in that region. Finally, we perform a pivoting step that yields a new, feasible region $R_{\mathbf{t}'}$ such that the line segment $\Pi_{\mathbf{t}'}$ in this new region extends the line segment $\Pi_{\mathbf{t}}$ of the old section. Then we iterate these steps until we end up in a region where we can extend the solution curve to infinity. Algorithm 1 summarizes these steps in pseudo-code. We proceed by describing the steps of the algorithm in more detail.

Algorithm 1 (Basic Algorithm)

Input: Directed graph G with piecewise linear, homogeneous, strictly increasing, and continuous marginal cost functions $f : \mathbb{R} \to \mathbb{R}$, demand vector $\mathbf{q}$
Output: Piecewise function $\mathbf{x} : \mathbb{R}_{\geq 0} \to \mathbb{R}^m$ s.t. $\mathbf{x}(\lambda)$ is a mincost flow for all $\lambda \geq 0$
choose $\boldsymbol{t}$ such that $\mathbf{0} \in R_{\mathbf{t}}$;
compute $\mathbf{L}_{\mathbf{t}}^*$;
$\lambda^{\min} \leftarrow 0$;
repeat
  compute $\boldsymbol{\pi}_{\mathbf{t}}, \Delta\boldsymbol{\pi}_{\mathbf{t}}$ and $\lambda_{\mathbf{t}}^{\max}$;
  $\mathbf{x}(\lambda) \leftarrow \mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^\top(\pi_{\mathbf{t}} + \lambda\Delta\boldsymbol{\pi}_{\mathbf{t}}) - \mathbf{d}_{\mathbf{t}}$ for $\lambda \in [\lambda_{\mathbf{t}}^{\min}, \lambda_{\mathbf{t}}^{\max}]$;$\lambda^{\min} \leftarrow \lambda_{\mathbf{t}}^{\max}$;
  $(\mathbf{t}, \mathbf{L}_{\mathbf{t}}^*) \leftarrow$ `PivotStep`$(R_{\mathbf{t}})$;
until $\lambda_{\mathbf{t}}^{\max} = \infty$;
return $\boldsymbol{x}$

3.3.1. Initialization. For $\lambda = 0$, the electrical flow is the trivial zero flow $\mathbf{x} = \mathbf{0}$ with potential $\boldsymbol{\pi} = \mathbf{0}$. We find the vector $\mathbf{t}^0$ such that $x_e = 0 \in [\tau_{e,t_e^0}, \tau_{e,t_e^0+1}]$ is satisfied for all $e \in E$ by iterating through all edges and all breakpoints of the cost functions, and we obtain the first region $R_{\mathbf{t}^0}$. For this region, we set up the Laplacian matrix $\mathbf{L}_{\mathbf{t}^0}$ and compute its generalized inverse $\mathbf{L}_{\mathbf{t}^0}^*$.

3.3.2. Main Loop. In every iteration of the main loop, we assume that we are given some region $R_{\mathbf{t}^i}$ and the generalized inverse of the corresponding Laplacian matrix $\mathbf{L}_{\mathbf{t}^i}^*$. The latter enables us to compute the potential offset and direction vectors $\boldsymbol{\pi}_{\mathbf{t}^i}$ and $\Delta\boldsymbol{\pi}_{\mathbf{t}^i}$ of the current region. With these vectors, we obtain the values $\underline{\lambda}_{\mathbf{t}^i}(e)$ and $\overline{\lambda}_{\mathbf{t}^i}(e)$ for all edges $e \in E$ and thus the interval $[\lambda_{\mathbf{t}^i}^{\min}, \lambda_{\mathbf{t}^i}^{\max}]$.

To find the next region $R_{\mathbf{t}^{i+1}}$, we introduce the notion of neighboring regions. We call two regions $R_{\mathbf{t}^1}, R_{\mathbf{t}^2}$ *neighboring* if $\mathbf{t}^2 = \mathbf{t}^1 \pm \mathbf{u}_e$ for some edge e, where $\mathbf{u}_e$ is the unit vector with the nonzero entry in the component corresponding to the edge e. The edge e is called *boundary edge* in this case. Two neighboring regions are separated by a hyperplane induced by the equation $\boldsymbol{\gamma}_e^\top \boldsymbol{\pi} = \sigma_{e,t_e}$, where e is the boundary edge. Let

$$E_{\mathbf{t}^i}^* := \{e \in E \mid \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i} < 0 \text{ and } \underline{\lambda}_{\mathbf{t}^i}(e) = \lambda_{\mathbf{t}^i}^{\max} \text{ or } \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i} > 0 \text{ and } \overline{\lambda}_{\mathbf{t}^i}(e) = \lambda_{\mathbf{t}^i}^{\max}\}$$

be the set of all edges that contribute to the minimum in the computation of $\lambda_{\mathbf{t}^i}^{\max}$. We will show that the neighboring regions of $R_{\mathbf{t}^i}$ that are induced by boundary edges $e^* \in E_{\mathbf{t}^i}^*$ are candidates for the next region. There are three possible cases:

Figure 4. Two possible pivoting situations: (a) $|E^*| = 1$: There is a unique boundary edge e^* inducing exactly one hyperplane to cross. The potential directions $\Delta\pi_{\mathbf{t}^i}$ and $\Delta\pi_{\mathbf{t}^{i+1}}$ have the same orientation wrt. the vector $\boldsymbol{\gamma}_{e^*}$ as shown in Theorem 2(ii). (b) $|E^*| > 1$: Multiple hyperplanes intersect in the boundary point of $R_{\mathbf{t}^i}$. The solution curve continues in exactly one of the adjacent regions.

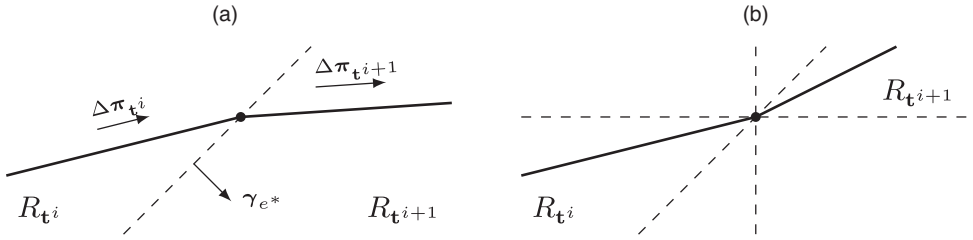

Non-degenerate pivoting step. Degenerate pivoting step.

1. There is no boundary edge, that is, $E^*_{\mathbf{t}^i} = \emptyset$. Then $\lambda^{\max}_{\mathbf{t}^i} = \infty$ and we stop.
2. There is a unique boundary edge $e^* \in E^*_{\mathbf{t}^i}$ inducing $\lambda^{\max}_{\mathbf{t}^i}$ i.e. $|E^*_{\mathbf{t}^i}| = 1$.
3. There are multiple edges inducing $\lambda^{\max}_{\mathbf{t}^i}$, that is, $|E^*_{\mathbf{t}^i}| > 1$. We say the region $R_{\mathbf{t}^i}$ is *degenerate*.

In general, we say an instance of Parametric Electrical Flow is a *degenerate instance* if there are feasible, degenerate regions. Figure 4 depicts the two possible pivoting situation: In Figure 4(a) there is a unique boundary that is crossed by the solution curve while in Figure 4(b) multiple boundaries intersect in one point traversed by the solution curve leading to a degenerate pivoting step.

3.3.3. Pivoting. First, we establish the following theorem about neighboring regions that proves a connection between the (generalized inverses of the) Laplacian matrices of two neighboring regions. In the following theorem and throughout the paper, $\operatorname{sgn} : \mathbb{R} \to \mathbb{R}$ denotes the sign function defined as $\operatorname{sgn}(x) = -1$ if $x < 0$, $\operatorname{sgn}(x) = 1$ if $x > 0$, and $\operatorname{sgn}(x) = 0$ otherwise.

Theorem 2. *Let $R_{\mathbf{t}^i}$ and $R_{\mathbf{t}^{i+1}}$ be two neighboring regions with the boundary edge e and let $\Delta c_e := c_{e,t_e^{i+1}} - c_{e,t_e^i}$. Then,*

i. *The generalized inverse of the Laplacian matrix in region $R_{\mathbf{t}^{i+1}}$ mapping to Π is*

$$\mathbf{L}^*_{\mathbf{t}^{i+1}} = \left(\mathbf{I}_n - \frac{\Delta c_e}{1 + \Delta c_e \boldsymbol{\gamma}_e^\top \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e} \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e \boldsymbol{\gamma}_e^\top\right) \mathbf{L}^*_{\mathbf{t}^i}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix.

ii. *The potential directions of the neighboring regions satisfy* $\operatorname{sgn}(\boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i}) = \operatorname{sgn}(\boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^{i+1}})$.

Proof. By the definition of Δc_e, we have $\mathbf{C}_{\mathbf{t}^{i+1}} = \mathbf{C}_{\mathbf{t}^i} + \operatorname{diag}(0, \ldots, 0, \Delta c_e, 0, \ldots, 0)$. Therefore, we have $\mathbf{L}_{\mathbf{t}^{i+1}} = \boldsymbol{\Gamma}\mathbf{C}_{\mathbf{t}^{i+1}}\boldsymbol{\Gamma}^\top = \mathbf{L}_{\mathbf{t}^i} + \Delta c_e \boldsymbol{\gamma}_e \boldsymbol{\gamma}_e^\top$. Denoting by $\hat{\boldsymbol{\gamma}}$ the vector obtained from $\boldsymbol{\gamma}$ by deleting the first component, we then obtain

$$\hat{\mathbf{L}}_{\mathbf{t}^{i+1}} = \hat{\mathbf{L}}_{\mathbf{t}^i} + \Delta c_e \hat{\boldsymbol{\gamma}}_e \hat{\boldsymbol{\gamma}}_e^\top. \tag{12}$$

The Sherman-Morrison formula (Hager [25], Sherman and Morrison [53]) yields

$$\hat{\mathbf{L}}^{-1}_{\mathbf{t}^{i+1}} = \left(\mathbf{I}_{n-1} - \frac{\Delta c_e}{1 + \Delta c_e \hat{\boldsymbol{\gamma}}_e^\top \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i} \hat{\boldsymbol{\gamma}}_e} \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i} \hat{\boldsymbol{\gamma}}_e \hat{\boldsymbol{\gamma}}_e^\top\right) \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i}.$$

The definition of the generalized inverse $\mathbf{L}^*$ implies that $\hat{\boldsymbol{\gamma}}_e^\top \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i} \hat{\boldsymbol{\gamma}}_e = \boldsymbol{\gamma}_e^\top \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e$. Because $\mathbf{L}^*$ is $\hat{\mathbf{L}}^{-1}$ with an additional zero row and column, we directly obtain (i).

For any two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, the determinant identity $\det(\mathbf{I}_n + \mathbf{v}\mathbf{w}^\top) = 1 + \mathbf{w}^\top \mathbf{v}$ holds true (Harville [27, corollary 18.1.3]). Hence, we obtain

$$\det(\hat{\mathbf{L}}_{\mathbf{t}^{i+1}}) = \det\left(\mathbf{I}_{n-1} + \Delta c_e \hat{\boldsymbol{\gamma}}_e \hat{\boldsymbol{\gamma}}_e^\top \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i}\right) \det\left(\hat{\mathbf{L}}_{\mathbf{t}^i}\right) = \left(1 + \Delta c_e \hat{\boldsymbol{\gamma}}_e^\top \hat{\mathbf{L}}^{-1}_{\mathbf{t}^i} \hat{\boldsymbol{\gamma}}_e\right) \det(\hat{\mathbf{L}}_{\mathbf{t}^i})$$

and thus

$$\begin{aligned}
\boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}^{i+1} &= \boldsymbol{\gamma}_e^\top \left(\mathbf{I}_n - \frac{\Delta c_e}{1 + \Delta c_e \boldsymbol{\gamma}_e^\top \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e} \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e \boldsymbol{\gamma}_e^\top\right) \mathbf{L}^*_{\mathbf{t}^i} \mathbf{q} \\
&= \frac{1}{1 + \Delta c_e \boldsymbol{\gamma}_e^\top \mathbf{L}^*_{\mathbf{t}^i} \boldsymbol{\gamma}_e} \boldsymbol{\gamma}_e^\top \mathbf{L}^*_{\mathbf{t}^i} \mathbf{q} = \frac{\det(\hat{\mathbf{L}}_{\mathbf{t}^{i+1}})}{\det(\hat{\mathbf{L}}_{\mathbf{t}^i})} \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}^i.
\end{aligned}$$

The matrices $\hat{\mathbf{L}}_{\mathbf{t}^i}$ and $\hat{\mathbf{L}}_{\mathbf{t}^i+1}$ are positive definite by Lemma 3(iv) implying $\det(\hat{\mathbf{L}}_{\mathbf{t}^i})/\det(\hat{\mathbf{L}}_{\mathbf{t}^{i+1}}) > 0$ and proving (ii). □

Theorem 2(i) provides an easy update formula for the generalized inverse of the neighboring region. In fact, there is just a simple rank 1 update needed to obtain the new generalized inverse. Using part (ii) of Theorem 2, we obtain the following corollary.

Corollary 2. *Let $R_{\mathbf{t}}$ be a feasible region and $e \in E^*_{\mathbf{t}}$. Furthermore, let $\mathbf{t}' := \mathbf{t} + \operatorname{sgn}(\boldsymbol{\gamma}_e^\top \Delta\pi_{\mathbf{t}})\mathbf{u}_e$. Then, $R_{\mathbf{t}'}$ is feasible and $\lambda^{\min}_{\mathbf{t}'} = \lambda^{\max}_{\mathbf{t}}$.*

Proof. Because $e \in E^*_{\mathbf{t}}$, the potential $\boldsymbol{\pi}_{\mathbf{t}} + \lambda^{\max}_{\mathbf{t}} \Delta\boldsymbol{\pi}_{\mathbf{t}}$ lies on the boundary between $R_{\mathbf{t}}$ and $R_{\mathbf{t}'}$. Thus, $\boldsymbol{\pi}_{\mathbf{t}} + \lambda^{\max}_{\mathbf{t}} \Delta\boldsymbol{\pi}_{\mathbf{t}} \in R_{\mathbf{t}'}$. Because all flows and potentials are unique, this potential is still a potential for an electrical flow for demand $\lambda\mathbf{q}$. Thus, $\lambda^{\max}_{\mathbf{t}} \geq \lambda^{\min}_{\mathbf{t}'}$. Assume that $\boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}} > 0$. Then $\lambda^{\max}_{\mathbf{t}}$ was induced by the breakpoint σ_{e,t_e+1}, that is, $\lambda^{\max}_{\mathbf{t}} = \overline{\lambda}_{\mathbf{t}}(e)$. By definition of $\mathbf{t}'$, $\sigma_{e,t_e+1} = \sigma_{e,t'_e}$. Furthermore, by Theorem 2(ii), we have that $\boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}'} > 0$. Thus, the value $\underline{\lambda}_{\mathbf{t}'}(e)$ contributes to the maximum in the computation of $\lambda^{\min}_{\mathbf{t}'}$. Because the potential $\boldsymbol{\pi}_{\mathbf{t}} + \lambda^{\max}_{\mathbf{t}} \Delta\boldsymbol{\pi}_{\mathbf{t}}$ lies on the boundary between the two regions, $\underline{\lambda}_{\mathbf{t}'} = \lambda^{\max}_{\mathbf{t}}$, and thus $\lambda^{\max}_{\mathbf{t}} \geq \lambda^{\min}_{\mathbf{t}'}$ and the claim follows. □

3.3.3.1. Pivoting in Nondegenerate Regions.

Assume that the current region $R_{\mathbf{t}}$ is nondegenerate, that is, $|E^*_{\mathbf{t}}| = 1$. Then there is a unique $e^* \in E$ such that $E^*_{\mathbf{t}} = \{e^*\}$. We define the *nondegenerate pivot rule* as

$$\mathbf{t} \longmapsto \mathbf{t}' := \mathbf{t} + \operatorname{sgn}(\boldsymbol{\gamma}_{e^*}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}})\mathbf{u}_{e^*}. \tag{13}$$

Corollary 2 implies that the pivoting rule from (13) yields a feasible region with $\lambda^{\min}_{\mathbf{t}'} = \lambda^{\max}_{\mathbf{t}}$, and hence the algorithm can proceed in this region.

3.3.3.2. Pivoting in Degenerate Regions.

We want to derive a pivoting rule that is well defined in every region, even in the presence of degenerate regions. The basic idea for this is the following. In addition to the solution curve $\lambda \longmapsto \boldsymbol{\pi}(\lambda)$ we define a new solution curve $\lambda \longmapsto \boldsymbol{\pi}^\varepsilon(\lambda)$ that is defined like the original solution curve but starts in a slightly perturbed initial potential $\boldsymbol{\pi}^0 + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ is a small vector depending on a fixed $\varepsilon > 0$. Then, we prove that $\boldsymbol{\pi}(\cdot)$ and $\boldsymbol{\pi}^\varepsilon(\cdot)$ have the same feasible regions, except for some regions that are not relevant for the solution. We further show that the curve $\boldsymbol{\pi}^\varepsilon(\cdot)$ has no degenerate regions. Hence, we can compute it with the unique pivoting rule for nondegenerate regions stated before. See Figure 5(a) for a visual respresentation of this idea. Because both curves have the same feasible regions, we derive a unique pivoting rule for the original solution curve from the behavior of the perturbed curve $\boldsymbol{\pi}^\varepsilon(\cdot)$, even in case of degeneracy. In particular, we will argue that it is not necessary to compute the perturbed curve explicitly, but that it is possible to retrace its course implicitly, yielding a lexicographic pivoting rule.

Formally, for some $\varepsilon > 0$, we define the potential perturbation vector $\boldsymbol{\varepsilon} := (0, \varepsilon^1, \varepsilon^2, \ldots, \varepsilon^{n-1}) \in \Pi$. Analogous to the initial potential of the original solution curve $\boldsymbol{\pi}^{\mathbf{0}} = \mathbf{0}$, we define the initial potential of the perturbed solution curve $\boldsymbol{\pi}^{0,\varepsilon} := \boldsymbol{\pi}^0 + \boldsymbol{\varepsilon}$.

First, we show that the initial region of the perturbed solution curve is well defined and that we can find it polynomial time. See also Figure 5(b) that depicts the relation between the potential $\boldsymbol{\pi}^{\mathbf{0}}$ and the initial region $R_{\mathbf{t}^{0,\varepsilon}}$.

Lemma 4. *There is a unique region $R_{\mathbf{t}^{0,\varepsilon}}$ and $\varepsilon^* > 0$ such that $\boldsymbol{\pi}^0 \in R_{\mathbf{t}^{0,\varepsilon}}$ and $\boldsymbol{\pi}^0 + \boldsymbol{\varepsilon} \in R_{\mathbf{t}^{0,\varepsilon}}$ for all $0 < \varepsilon < \varepsilon^*$. Furthermore, the region $R_{\mathbf{t}^{0,\varepsilon}}$ can be found in $\mathcal{O}(m)$ time.*

Proof. First, observe that for every edge $e \in E$, there are numbers $j_1, j_2 \in \mathbb{N}$ such that $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} = \varepsilon^{j_1} - \varepsilon^{j_2}$ or, if e is incident to the vertex with index 1, $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} = \pm\, \varepsilon^{j_1}$, that is, $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}$ is either a monomial in ε or the difference of two such monomials. Thus, all roots of $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}$ can only be zero or one and we get that, for all $0 < \varepsilon < 1$, the sign of $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}$ is constant and $|\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}| < \varepsilon$. We claim that there is a unique region $R_{\mathbf{t}}$ satisfying both

$$\boldsymbol{\pi}^0 \in R_{\mathbf{t}}, \quad \text{and} \quad \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 > \sigma_{e,t_e} \quad \text{for all edges } e \in E. \tag{14}$$

To prove the claim, we consider the following construction. Choose any region containing $\boldsymbol{\pi}^0$. If $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} = \sigma_{e,t_e}$ for some edge, then move to the adjacent region with $t'_e = t_e - 1$. This ensures that $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} < \sigma_{e,t'_e}$ without violating the feasibility of $\boldsymbol{\pi}^0$.

We partition the edges into sets

$$\begin{aligned}
\bar{E}_1 &:= \{e \in E \mid \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 < \sigma_{e,t_e+1}\},\\
\bar{E}_2 &:= \{e \in E \mid \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 = \sigma_{e,t_e+1} \text{ and } \operatorname{sgn}(\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}) = 1\},\\
\bar{E}_3 &:= \{e \in E \mid \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 = \sigma_{e,t_e+1} \text{ and } \operatorname{sgn}(\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}) = -1\}.
\end{aligned}$$

Figure 5. (a) The solution curve (thick line) with a nondegenerate and a degenerate pivot step. The perturbed solution curve (dotted line) starts in the initial potential $\boldsymbol{\pi}^{0,\varepsilon}$ perturbed by the vector $\boldsymbol{\varepsilon}$. It performs the same nondegenerate boundary crossing, whereas there are several, nondegenerate boundary crossings for the degenerate crossing of the original solution curve. (b) The initial potential $\boldsymbol{\pi}^0$ is feasible in multiple regions. The dotted curve shows all possible perturbed initial potentials $\boldsymbol{\pi}^{0,\varepsilon} = \boldsymbol{\pi}^0 + \boldsymbol{\varepsilon}$ for all $\varepsilon > 0$. The highlighted regions mark the regions $R_{\mathbf{t}}$ and $R_{\mathbf{t}^{0,\varepsilon}}$ defined in Lemma 4. For small ε, the curve of all perturbed potentials $\boldsymbol{\pi}^{0,\varepsilon} = \boldsymbol{\pi}^0 + \boldsymbol{\varepsilon}$ is in the fixed, unique region $R_{\mathbf{t}^{0,\varepsilon}}$.

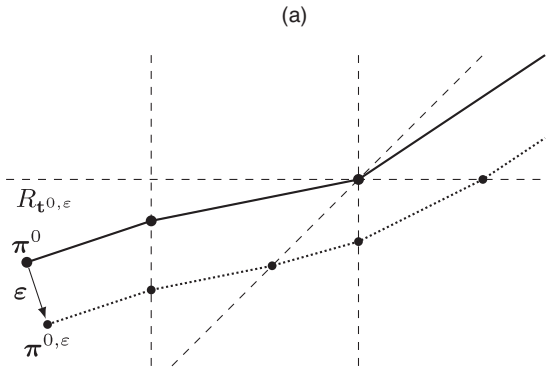

Solution curve and perturbed solution curve.

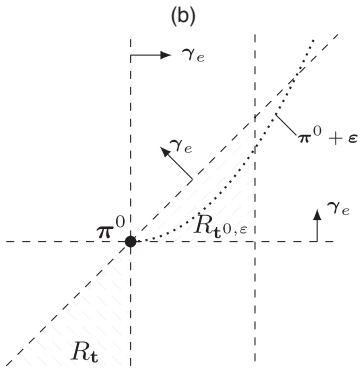

Start region $R_{\mathbf{t}^{\varepsilon}}$.

We define

$$\tilde{\varepsilon}_1 := \min\{\sigma_{e,t_e+1} - \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 \mid e \in \bar{E}_1\},$$
$$\tilde{\varepsilon}_2 := \min\{\sigma_{e,t_e+2} - \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 \mid e \in \bar{E}_2\},$$
$$\tilde{\varepsilon}_3 := \min\{\boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 - \sigma_{e,t_e} \mid e \in E\},$$
$$\varepsilon^* := \min\{\tilde{\varepsilon}_1, \tilde{\varepsilon}_2, \tilde{\varepsilon}_3, 1\}.$$

Note that $\tilde{\varepsilon}_1, \tilde{\varepsilon}_2, \tilde{\varepsilon}_3 > 0$ and thus $\varepsilon^* > 0$.

We claim that $R_{\mathbf{t}^{0,\varepsilon}}$ with $\mathbf{t}^{0,\varepsilon} := \mathbf{t} + \sum_{e \in \bar{E}_2} \mathbf{u}_e$ has the claimed properties. Using (14), we have that $\sigma_{e,t_e} \le \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}^0 \le \sigma_{e,t_e+1}$ for every edge $e \in E$. Because $t_e^\varepsilon = t_e$ for all edges $e \notin E_2$, we get by the definition of ε^* and the fact that $|\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}| < \varepsilon$ that

$$\sigma_{e,t_e^\varepsilon} = \sigma_{e,t_e} < \boldsymbol{\pi}^0 + \boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} < \sigma_{e,t_e+1} = \sigma_{e,t_e^\varepsilon+1},$$

where the last inequality follows again by the definition of ε^* for all edges $e \in \bar{E}_1$, and by the fact that $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} < 0$ for all edges $e \in \bar{E}_3$. For $e \in \bar{E}_2$, we get

$$\sigma_{e,t_e^\varepsilon} = \sigma_{e,t_e+1} = \boldsymbol{\pi}^0 < \boldsymbol{\pi}^0 + \boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} < \sigma_{e,t_e+2} = \sigma_{e,t_e^\varepsilon+1}$$

by $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} > 0$ and the definition of $\tilde{\varepsilon}_3$. Hence, $\boldsymbol{\pi}^0 + \boldsymbol{\varepsilon} \in R_{\mathbf{t}^{0,\varepsilon}}$, as well as $\boldsymbol{\pi}^0 \in R_{\mathbf{t}^{0,\varepsilon}}$ by definition.

The region $R_{\mathbf{t}^\varepsilon}$ can be found efficiently in the following steps. (1) Find a region that contains $\boldsymbol{\pi}^0$ by checking all breakpoints for all edges. (2) Find the region $R_{\mathbf{t}}$ as previously claimed by checking if $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon} = \sigma_{e,t_e}$. (3) Find the region $R_{\mathbf{t}^{0,\varepsilon}}$ by checking whether $\operatorname{sgn}(\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon})$ is positive or negative for all $e \in \bar{E}_2$. By the previous considerations, this can be checked by comparing the exponents j_1 and j_2 of the polynomial $\boldsymbol{\gamma}_e^\top \boldsymbol{\varepsilon}$. □

Let $R_{\mathbf{t}^{0,\varepsilon}}$ be the region defined in Lemma 4. Then we define the *perturbed solution curve* $\lambda \longmapsto \boldsymbol{\pi}^\varepsilon(\lambda)$ as follows: For every $\lambda \ge 0$, the potential $\boldsymbol{\pi}^\varepsilon(\lambda)$ induces an equilibrium flow with the excess $\lambda \mathbf{q} + \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon}$. Recall, that the solution curve $\boldsymbol{\pi}(\cdot)$ is defined likewise but induces flows with excess vectors $\lambda \mathbf{q}$. Thus, we can interpret the perturbed solution curve as a solution to our initial problem, but with a small perturbation on the excess vectors; that is, every electrical flow induced by the perturbed solution curve is perturbed by a small auxiliary flow.

The perturbed solution curve is computed like the original solution curve, but the Laplacian system $\mathbf{L}_{\mathbf{t}} \boldsymbol{\pi} - \tilde{\mathbf{d}}_{\mathbf{t}} = \lambda \mathbf{q}$ is replaced by $\mathbf{L}_{\mathbf{t}} \boldsymbol{\pi} - \tilde{\mathbf{d}}_{\mathbf{t}} = \lambda \mathbf{q} + \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon}$. Then, we characterize the segments of the perturbed solution curve in region $R_{\mathbf{t}}$ (i.e., potentials in region $R_{\mathbf{t}}$ that induce flows with the correct excess) as before by

$$\Pi_{\mathbf{t}}^\varepsilon = \{\boldsymbol{\pi}_{\mathbf{t}} + \lambda \Delta \boldsymbol{\pi}_{\mathbf{t}} + \mathbf{L}_{\mathbf{t}}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon} \mid \lambda \in [\lambda_{\mathbf{t}}^{\min,\varepsilon}, \lambda_{\mathbf{t}}^{\max,\varepsilon}]\},$$

where we use the values

$$\underline{\lambda}_{\mathbf{t}}^\varepsilon(e) := \frac{\sigma_{e,t_e} - \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}_{\mathbf{t}} - \boldsymbol{\gamma}_e^\top \mathbf{L}_{\mathbf{t}}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon}}{\boldsymbol{\gamma}_e^\top \Delta \pi_{\mathbf{t}}} \quad \text{and} \quad \overline{\lambda}_{\mathbf{t}}^\varepsilon(e) := \frac{\sigma_{e,t_e+1} - \boldsymbol{\gamma}_e^\top \boldsymbol{\pi}_{\mathbf{t}} - \boldsymbol{\gamma}_e^\top \mathbf{L}_{\mathbf{t}}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon}}{\boldsymbol{\gamma}_e^\top \Delta \pi_{\mathbf{t}}}$$

and define $\lambda_{\mathbf{t}}^{\min,\varepsilon}$ and $\lambda_{\mathbf{t}}^{\max,\varepsilon}$ analogously to the unperturbed case. We define for every edge $e \in E$ the vector $\mathbf{m}_{e,\mathbf{t}} := -\frac{1}{\boldsymbol{\gamma}_e^\top \Delta\pi_{\mathbf{t}}} \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \mathbf{L}_{\mathbf{t}}^* \boldsymbol{\gamma}_e$ and the polynomial

$$p_{e,\mathbf{t}}(\varepsilon) := \mathbf{m}_{e,\mathbf{t}}^\top \boldsymbol{\varepsilon} = -\frac{1}{\boldsymbol{\gamma}_e^\top \Delta\pi_{\mathbf{t}}} \boldsymbol{\gamma}_e^\top \mathbf{L}_{\mathbf{t}}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \boldsymbol{\varepsilon}.$$

Then we observe that $\underline{\lambda}_{\mathbf{t}}^{\varepsilon}(e) = \underline{\lambda}_{\mathbf{t}}(e) + p_{e,\mathbf{t}}(\varepsilon)$ and $\overline{\lambda}_{\mathbf{t}}^{\varepsilon}(e) = \overline{\lambda}_{\mathbf{t}}(e) + p_{e,\mathbf{t}}(\varepsilon)$. We say a vector $\boldsymbol{v}^1 = (v_1^1, \dots, v_n^1) \in \mathbb{R}^n$ is lexicographically smaller than the vector $\boldsymbol{v}^2 = (v_1^2, \dots, v_n^2) \in \mathbb{R}^n$ if there is $i \in \{2, \dots, n\}$[1] with $v_j^1 = v_j^2$ for all $j < i$ and $v_i^1 < v_i^2$ and denote this by $\boldsymbol{v}^1 \prec_{\text{LEX}} \boldsymbol{v}^2$. The following lemma states some basic properties of the perturbation $p_{e,\mathbf{t}}(\varepsilon)$ added to the values $\underline{\lambda}_{\mathbf{t}}^{\varepsilon}(e)$ and $\overline{\lambda}_{\mathbf{t}}^{\varepsilon}(e)$ of the perturbed solution curve.

Lemma 5. *For any region $R_{\mathbf{t}}$ let $\mathbf{X}_{\mathbf{t}} := \mathbf{L}_{\mathbf{t}}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}}$. Then,*
i. *the left nullspace of $\mathbf{X}_{\mathbf{t}}$ is $\ker(\mathbf{X}_{\mathbf{t}}^\top) = \operatorname{span}(\mathbf{u}_1)$, where $\mathbf{u}_1 \in \mathbb{R}$ is the first unit vector.*
ii. *We have $\mathbf{X}_{\mathbf{t}}^\top \boldsymbol{\Gamma} \mathbf{v} = \mathbf{0}$ if and only if $\mathbf{v} = \mathbf{0} \in \mathbb{R}^m$.*
iii. *There is $\varepsilon_{\mathbf{t}}^*$ such that for all $0 < \varepsilon < \varepsilon_{\mathbf{t}}^*$ and any two edges $e_1, e_2 \in E$, we have $p_{e_1,\mathbf{t}}(\varepsilon) < p_{e_2,\mathbf{t}}(\varepsilon)$ if and only if $\mathbf{m}_{e_1,\mathbf{t}} \prec_{\text{LEX}} \mathbf{m}_{e_2,\mathbf{t}}$.*

Proof. By definition of $\mathbf{L}_{\mathbf{t}}^*$ we get that $\mathbf{X}_{\mathbf{t}}^\top \mathbf{u}_1 = \mathbf{L}_{\mathbf{t}^\varepsilon} \mathbf{L}_{\mathbf{t}}^* \mathbf{u}_1 = 0$. The matrix $\mathbf{L}_{\mathbf{t}^{0,\varepsilon}}$ is invertible on Π; thus, by the properties of generalized inverses, $\mathbf{L}_{\mathbf{t}^\varepsilon}^* \mathbf{L}_{\mathbf{t}^{0,\varepsilon}}$ is the projection onto Π. Because $\mathbf{L}_{\mathbf{t}}^*$ maps to Π, we obtain $\mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \mathbf{L}_{\mathbf{t}}^* \mathbf{v} = 0$ only if $\mathbf{L}_{\mathbf{t}}^* \mathbf{v} = 0$. Because the nullspace of $\mathbf{L}_{\mathbf{t}}^*$ is the span of $\mathbf{u}_1$, (i) follows. Every column of $\boldsymbol{\Gamma}$ has zero sum; hence, $\boldsymbol{\Gamma}\mathbf{v}$ has also zero sum for any $\mathbf{v} \in \mathbb{R}^m$. Thus, $\mathbf{u}_1$ is not in the column space of $\boldsymbol{\Gamma}$, and (ii) follows by (i). Finally, (iii) follows by the fact, that $p_{e,\mathbf{t}}$ is a polynomial in ε with $p_{e,\mathbf{t}}(0) = 0$ with coefficient vector $\mathbf{m}_{e,\mathbf{t}}$. □

We proceed to prove a result relating the feasible regions of the perturbed and the original solution curve.

Lemma 6. *There exists $\varepsilon^* > 0$ such that, for all $0 < \varepsilon < \varepsilon^*$, the following holds.*
i. *Let $R_{\mathbf{t}}$ be an infeasible region, that is, $\Pi_{\mathbf{t}} = \emptyset$. Then $\Pi_{\mathbf{t}}^{\varepsilon} = \emptyset$ as well.*
ii. *Let $R_{\mathbf{t}}$ be a feasible region with $\lambda_{\mathbf{t}}^{\max} > \lambda_{\mathbf{t}}^{\min}$. Then $\lambda_{\mathbf{t}}^{\max,\varepsilon} > \lambda_{\mathbf{t}}^{\min,\varepsilon}$.*
iii. *Let $R_{\mathbf{t}}$ be a nondegenerate region, that is, $|E_{\mathbf{t}}^*| \le 1$. Then, $E_{\mathbf{t}}^{*,\varepsilon} = E_{\mathbf{t}}^*$.*

Proof. Because $p_{e,\mathbf{t}}(\varepsilon)$ is a polynomial in ϵ with $p_{e,\mathbf{t}}(0) = 0$, we can find, for every $\beta > 0$, an $\varepsilon^* > 0$ such that $|p_{e,\mathbf{t}}(\varepsilon)| < \beta$ for all $0 < \varepsilon < \varepsilon^*$.

Consider some fixed, infeasible region $R_{\mathbf{t}}$. Then, $\lambda_{\mathbf{t}}^{\min} > \lambda_{\mathbf{t}}^{\max}$. If one of these values is infinite (because the maximum/minimum was taken over an empty set), then the corresponding values $\lambda_{\mathbf{t}}^{\min,\varepsilon}$ or $\lambda_{\mathbf{t}}^{\max,\varepsilon}$ are also infinite and thus $\Pi_{\mathbf{t}}^{\varepsilon} = \emptyset$ as well. Now, assume that $\lambda_{\mathbf{t}}^{\min} > \lambda_{\mathbf{t}}^{\max}$ and both values are finite. Then, $\beta := \lambda_{\mathbf{t}}^{\min} - \lambda_{\mathbf{t}}^{\max} > 0$. Choose $\varepsilon_{\mathbf{t}}^{*,1}$ such that $|p_{e,\mathbf{t}}(\varepsilon)| < \frac{\beta}{2}$ for all edges $e \in E$ and all $0 < \varepsilon < \varepsilon_{\mathbf{t}}^{*,1}$. (This is possible because E is finite.) Denote by e_1 and e_2 the edges for which the maximum/minimum is attained in $\lambda_{\mathbf{t}}^{\min,\varepsilon}$ and $\lambda_{\mathbf{t}}^{\max,\varepsilon}$. Because $\underline{\lambda}_{\mathbf{t}}^{\varepsilon}(e) = \underline{\lambda}_{\mathbf{t}}(e) - p_{e,\mathbf{t}}(\varepsilon)$ and $\overline{\lambda}_{\mathbf{t}}^{\varepsilon}(e) = \overline{\lambda}_{\mathbf{t}}(e) - p_{e,\mathbf{t}}(\varepsilon)$, we get $\lambda_{\mathbf{t}}^{\min,\varepsilon} - \lambda_{\mathbf{t}}^{\max,\varepsilon} > \lambda_{\mathbf{t}}^{\min} - \lambda_{\mathbf{t}}^{\max} - p_{e_1,\mathbf{t}} + p_{e_2,\mathbf{t}} > \beta - \frac{\beta}{2} - \frac{\beta}{2} = 0$. Hence, for all $0 < \varepsilon < \varepsilon_{\mathbf{t}}^{\prime 1}$, the region $R_{\mathbf{t}}$ is infeasible in the degenerate instance as well. Because there is only a finite number of regions, $\varepsilon^{*,1} := \min\{\varepsilon_{\mathbf{t}}^{*,1} \mid \Pi_{\mathbf{t}} = \emptyset\} > 0$, and thus (i) is true when choosing $0 < \varepsilon^* \le \varepsilon^{*,1}$.

With the same argumentation (by setting $\beta := \lambda_{\mathbf{t}}^{\max} - \lambda_{\mathbf{t}}^{\min} > 0$), we find an $\varepsilon^{*,2}$, such that (ii) is true.

Finally, if $R_{\mathbf{t}}$ is nondegenerate, then there is a unique edge e^* inducing the minimum in the computation of $\lambda^{\max}$. Thus, similar to the first part of the proof, for every region $R_{\mathbf{t}}$, we can find an $\varepsilon_{\mathbf{t}}^{*,3}$ such that $|p_{e,\mathbf{t}}(\varepsilon)| < \frac{\beta}{2}$, where β is the difference of the second lowest value in the computation of the minimum for $\lambda_{\mathbf{t}}^{\max}$ and $\lambda_{\mathbf{t}}^{\max}$. Then the minimum of in the computation of $\lambda_{\mathbf{t}}^{\max,\varepsilon}$ is attained for the same (unique) edge e^* as in the minimization of $\lambda_{\mathbf{t}}^{\max}$, proving (iii). Using $\varepsilon^* := \min\{\varepsilon^{*,1}, \varepsilon^{*,2}, \varepsilon^{*,3}\}$ completes the proof. □

Lemma 6 shows that, for ε small enough, every region that is feasible for the perturbed solution curve is also feasible for the original solution curve. The converse is true for all regions with $\lambda_{\mathbf{t}}^{\max} > \lambda_{\mathbf{t}}^{\min}$ and, furthermore, for every nondegenerate region, the same edge induces the $\lambda^{\max}$ of the region for the perturbed and the unperturbed solution curve. Thus, the perturbed and the unperturbed solution curve behave similar in almost all regions. Only the regions with $\lambda_{\mathbf{t}}^{\max} = \lambda_{\mathbf{t}}^{\min}$ that are feasible for the solution curve can also be infeasible for the perturbed solution curve. However, these regions do not contribute to the actual solution curve, as they only define one point. Therefore, it is enough to track the regions that are feasible for the perturbed solution curve. We proceed to show that we can compute the course of the perturbed solution curve with our basic algorithm as there are no degenerate regions for the perturbed solution curve.

Theorem 3. *There is $\varepsilon^* > 0$ such that for all $0 < \varepsilon < \varepsilon^*$ and every region $R_{\mathbf{t}}$, we have $|E_{\mathbf{t}}^{*,\varepsilon}| \le 1$.*

Proof. Assume toward a contradiction that in every interval $(0, \varepsilon^*)$, there are infinitely many values $\varepsilon_1 > \varepsilon_2 > \varepsilon_3 > \ldots > 0$ such that for every ε_i, there is a region $R_{\mathbf{t}}$ with $|E^{*,\varepsilon_i}_{\mathbf{t}}| \geq 2$ for all $i = 1, 2, \ldots$. Because there are only finitely many regions and also only finitely many (pairs of) edges, we can assume that there is one fixed region $R_{\mathbf{t}}$ and two distinct edges e_1, e_2 such that there are infinitely many values $\varepsilon_1 > \varepsilon_2 > \varepsilon_3 > \ldots > 0$ with $\{e_1, e_2\} \subseteq E^{*,\varepsilon_i}_{\mathbf{t}}$ for all $i = 1, 2, \ldots$. Without loss of generality, we furthermore assume that $\boldsymbol{\gamma}_{e_1}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}} > 0$ and $\boldsymbol{\gamma}_{e_2}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}} > 0$ because the orientation of the edges can be chosen arbitrarily. We then infer from the definition of $E^{*,\varepsilon_i}_{\mathbf{t}}$ that $\lambda^{\max,\varepsilon_i}_{\mathbf{t}} = \overline{\lambda}^{\varepsilon_i}_{\mathbf{t}}(e_1) = \overline{\lambda}^{\varepsilon_i}_{\mathbf{t}}(e_2)$. By choosing ε^* small enough, we ensure that this also implies that $\overline{\lambda}_{\mathbf{t}}(e_1) = \overline{\lambda}_{\mathbf{t}}(e_2)$. Hence, we know that, for all $i = 1, 2, \ldots$, we have $p_{e_1,\mathbf{t}}(\varepsilon_i) = p_{e_2,\mathbf{t}}(\varepsilon_i)$. Thus, the polynomial $p(\varepsilon) := p_{e_1,\mathbf{t}}(\varepsilon) - p_{e_2,\mathbf{t}}(\varepsilon)$ has infinitely many distinct roots and must therefore by the zero polynomial. By the definition of $p_{e,\mathbf{t}}(\varepsilon)$, this yields

$$\left(\frac{1}{\boldsymbol{\gamma}_{e_1}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}}} \boldsymbol{\gamma}_{e_1}^\top - \frac{1}{\boldsymbol{\gamma}_{e_2}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}}} \boldsymbol{\gamma}_{e_2}^\top\right) \mathbf{L}^*_{\mathbf{t}} \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} = 0.$$

By Lemma 5, this implies that $\boldsymbol{\gamma}_{e_1}^\top = \boldsymbol{\gamma}_{e_2}^\top$. Because we do not allow for parallel edges, we get $e_1 = e_2$, which is a contradiction. □

The previous results can be summarized as follows. Every feasible region of the perturbed solution curve is also a feasible region for the original solution curve. Furthermore, every nontrivial, feasible region of the original solution curve (that is a region $R_{\mathbf{t}}$ with $\lambda^{\min}_{\mathbf{t}} < \lambda^{\max}_{\mathbf{t}}$, i.e., a region that contributes a function part that defines more than a single point of the output functions $\mathbf{x}$) is also feasible for the perturbed solution curve. Hence, it is no restriction to only consider regions feasible for the perturbed solution curve. In nondegenerate regions, both solution curves have the same behavior (i.e., $\lambda^{\max}$ is determined by the same edge), so the pivot rule from the nondegenerate case can be used to track the behavior of the perturbed curve. In a degenerate region, we can find the unique edge in $E^{*,\varepsilon}_{\mathbf{t}}$ by finding the edge in $E^*_{\mathbf{t}}$ that minimizes $p_{e,\mathbf{t}}(\varepsilon)$ for small $\varepsilon > 0$. By Lemma 5, we know that this is equivalent to finding the lexicographic minimum of all vectors $\mathbf{m}_{e,\mathbf{t}}, e \in E^*_{\mathbf{t}}$. The function PivotStep gives the pseudocode of the previous procedure.

Function (PivotStep(**t**))
Input: Region vector **t** of some feasible region $R_{\mathbf{t}}$.
Output: A region vector of a neighboring region **t′** with $\lambda^{\max}_{\mathbf{t}} = \lambda^{\min}_{\mathbf{t}'}$.
$E^*_{\mathbf{t}} \leftarrow \{e \in E \mid \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i} < 0 \text{ and } \underline{\lambda}_{\mathbf{t}^i} = \lambda^{\max}_{\mathbf{t}^i} \text{ or } \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i} > 0 \text{ and } \overline{\lambda}_{\mathbf{t}^i} = \lambda^{\max}_{\mathbf{t}^i}\}$;
$\mathbf{m}_{e,\mathbf{t}} \leftarrow -\frac{1}{\boldsymbol{\gamma}_e^\top \Delta\pi_{\mathbf{t}}} \mathbf{L}_{\mathbf{t}^{0,\varepsilon}} \mathbf{L}^*_{\mathbf{t}} \boldsymbol{\gamma}_e$ for all $e \in E^*_{\mathbf{t}}$;
find $e^* \in E^*_{\mathbf{t}}$ with $\mathbf{m}_{e^*,\mathbf{t}} \prec_{\text{LEX}} \mathbf{m}_{e,\mathbf{t}}$ for all $e^* \neq e \in E^*_{\mathbf{t}}$;
$\mathbf{t}' \leftarrow \mathbf{t} + \operatorname{sgn}(\boldsymbol{\gamma}_{e^*}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}^i}) \mathbf{u}_{e^*}$;
$\mathbf{L}^*_{\mathbf{t}'} \leftarrow (\mathbf{I}_n - \frac{\Delta c_{e^*}}{1 + \Delta c_{e^*} \boldsymbol{\gamma}_{e^*}^\top \mathbf{L}^*_{\mathbf{t}} \boldsymbol{\gamma}_{e^*}} \mathbf{L}^*_{\mathbf{t}} \boldsymbol{\gamma}_{e^*} \boldsymbol{\gamma}_{e^*}^\top) \mathbf{L}^*_{\mathbf{t}}$;
return $(\mathbf{t}', \mathbf{L}^*_{\mathbf{t}'})$.

3.3.4. Termination of the Algorithm. By definition, the algorithm considers exactly one region $R_{\mathbf{t}}$ in every iteration. Hence, the number of iterations of the algorithm depends on the number of regions traversed by the solution curve.

Lemma 7. *The algorithm considers every region $R_{\mathbf{t}}$ at most once.*

Proof. In the nondegenerate case, the value $\lambda^{\min}_{\mathbf{t}^i}$ is strictly increasing in the course of the algorithm. Thus, the same region cannot be considered twice. This argument holds also for the perturbed instance; thus, using the lexicographic rule also implies that no region is considered twice. □

Because the number of iterations is bounded by the number of regions, the algorithm terminates within finite time if the graph and the number of breakpoints of the cost functions are finite. We discuss the computational complexity of the algorithm and the number of iterations performed by the algorithm in more detail in Section 6. In particular, we show that in general the latter can be exponential in the input size, that is, in the size of the network and the number of breakpoints of the functions f_e.

4. Electrical Flows with Discontinuous and Non-Homogeneous Marginal Costs

In this section we present several extensions to our basic electrical flow model and discuss how to adapt the basic algorithm to also work in these more general setting. In particular, we show how to use discontinuous marginal

cost functions to model directed edges and flow. This is the key to use the basic algorithm also in the setting of general mincost flows.

4.1. Discontinuous Marginal Costs

As a first generalization, we consider marginal cost functions $f_e : \mathbb{R} \to \mathbb{R}$ that are strictly increasing, piecewise linear, and homogeneous but *need not be continuous* (see Figure 6(a) for an example). We call a (possibly discontinuous) function homogeneous if $f_e^-(0) \le 0 \le f_e^+(0)$. We obtain the following optimality conditions as a corollary of Lemma 1.

Corollary 3. *A vector $\mathbf{x} \in \mathbb{R}^m$ is an electrical flow solving (6) if and only if there is a vector $\boldsymbol{\pi} \in \mathbb{R}^n$ of vertex potentials with*

$$f_e^-(x) \le \pi_w - \pi_v \le f_e^+(x) \quad \text{for all } e = (v,w) \in E, \tag{15a}$$

$$\boldsymbol{\Gamma}\mathbf{x} = \mathbf{q}. \tag{15b}$$

Proof. Setting the dual multiplier $\mu_e = 0$ for all $e \in E$ in the proof of Lemma 1 yields the claimed result. □

Despite the slightly different optimality conditions (15a), we still use the same approach as in Section 3.3. In order to compute an electrical, we find a potential $\boldsymbol{\pi}$ such that the (unique) flow, solving (15a) satisfies (15b) as well. In order to solve (15a), we need to define an inverse for every marginal cost function f_e. Let $\sigma_{e,t}^+ := f_e^+(\tau_{e,t_e})$ and $\sigma_{e,t}^- := f_e^-(\tau_{e,t_e})$ (Figure 6). Then, we define the (pseudo-)inverse of f_e by

$$f_e^{-1}(\alpha) = \begin{cases} \dfrac{1}{a_{e,t}}\alpha - \dfrac{b_{e,t}}{a_{e,t}} & \text{for all } \alpha \in (\sigma_{e,t}^+, \sigma_{e,t+1}^-) \\ \tau_{e,t} & \text{for all } \alpha \in [\sigma_{e,t}^-, \sigma_{e,t}^+]. \end{cases} \tag{16}$$

Given any potential $\boldsymbol{\pi}$, the flow $\mathbf{x} := \mathbf{f}^{-1}(\boldsymbol{\pi}) := (f_e^{-1}(\pi_w - \pi_v))_{e\in E}$ satisfies (15a) together with $\boldsymbol{\pi}$ by definition. For the sake of a simple notation, we refer to the breakpoints $\sigma_{e,t}^+, \sigma_{e,t}^-$ as $\sigma_{e,t}$ as we did before. With the explicit formula (16) for the inverse marginal cost function, we obtain as in Lemma 2 that the function $\mathbf{f}^{-1}$ has the explicit form

$$\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{C}_\mathbf{t}\boldsymbol{\Gamma}^\top\boldsymbol{\pi} - \mathbf{d}_\mathbf{t}$$

in the region $R_\mathbf{t}$. The only difference to the basic setting from Section 3.3 is that now, depending on the region, there may be one or more edges with constant inverse function f_e^{-1}. This means there might be some zero diagonal elements in the diagonal matrix $\mathbf{C}_\mathbf{t}$. In general, this is no problem as long as the Laplacian matrix of the region $\mathbf{L}_\mathbf{t} = \boldsymbol{\Gamma}\mathbf{C}_\mathbf{t}\boldsymbol{\Gamma}^\top$ has a unique inverse on Π. To formalize this, given a region $R_\mathbf{t}$, we call two vertices v, w *actively connected* if they are connected by an undirected path in G such that every edge in the path satisfies $c_{e,t_e} \neq 0$. An *actively connected compontent* is a maximal subset of vertices $V' \subseteq V$ where all vertices $v \in V'$ are actively connected. Lemma 3 implies that $\mathbf{L}_\mathbf{t}$ is has a unique inverse on Π as long as there is only one actively connected component.

We call a region where the Laplacian matrix $\mathbf{L}_\mathbf{t}$ has a nonunique inverse on Π, that is, where $\operatorname{rank}(\mathbf{L}_\mathbf{t}) < n-1$ *ambiguous region*. The name is motivated by the fact that, in these regions, there is no unique potential satisfying Condition (15a) for a given flow $\mathbf{x}$.

Unfortunately, we cannot guarantee that every region traversed by the solution curve is nonambiguous as we see in Example 3. However, we proceed to prove the following three properties of ambiguous regions:

Figure 6. Discontinuous marginal cost function (a) and its pseudo-inverse (b).

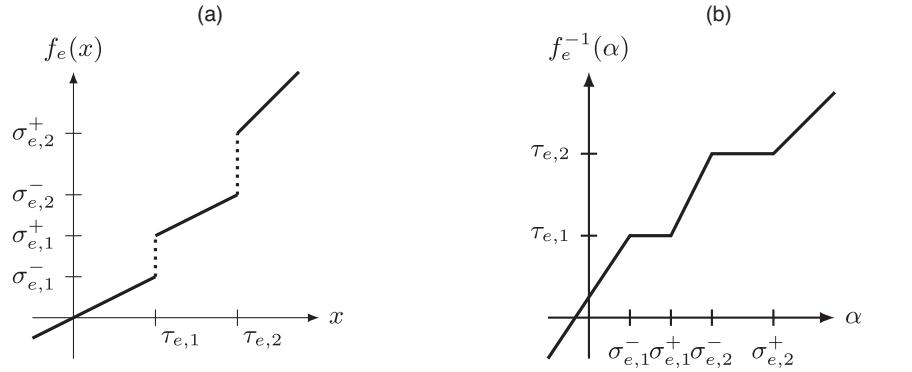

Marginal cost function function $f_e(x)$. Inverse marginal cost function $f_e^{-1}(\alpha)$.

i. Given a potential $\boldsymbol{\pi}$ that lies inside an ambiguous region $R_{\mathbf{t}}$, we can obtain a different potential $\boldsymbol{\pi}'$ in a different neighboring *nonambiguous region* $R_{\mathbf{t}'}$ that induces the same flow with respect to the optimality condition (15a), that is, via the inverse formula (16). Additionally, this potential and this region can be found in $\mathcal{O}(n^2)$ time.

ii. In the special case that the region $R_{\mathbf{t}}$ is not "too ambiguous," meaning that the Laplace matrix of the ambiguous regions has only one additional kernel dimension, there are only exactly two potentials in two different neighboring nonambiguous regions inducing the same flow. This is of particular interest because the algorithm will only encounter such regions as we will prove in the subsequent lemma.

iii. Furthermore, if we are in special case (ii) and are also already given one of these neighboring regions and its inverse Laplacian matrix, then we can find the other unique neighboring nonambiguous region and its inverse Laplacian matrix efficiently (namely, in $\mathcal{O}(n^2)$ time). This is actually the situation when the algorithm encounters an ambiguous region when pivoting to the next region—the ambiguous region can then be skipped, and the algorithm can proceed in the next, nonambiguous region without introducing any discontinuities in the flow functions. In particular, this skipping process has the same cost, $\mathcal{O}(n^2)$, just like a normal pivoting step.

These key properties are shown formally in the following theorem.

Theorem 4. *Let $\boldsymbol{\pi} \in R_{\mathbf{t}}$ be a potential in an ambiguous region $R_{\mathbf{t}}$. Denote by $K := n - 1 - \operatorname{rank}(\mathbf{L}_{\mathbf{t}})$ the dimension of the kernel of $\mathbf{L}_{\mathbf{t}}$ in Π. Then,*

i. *There is a potential $\boldsymbol{\pi}'$ and a nonambiguous region $R_{\mathbf{t}'}$ such that $\boldsymbol{\pi}' \in R_{\mathbf{t}'}$ and $\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\boldsymbol{\pi}')$. The potential and the region can be found in $\mathcal{O}(n^2)$ time.*

ii. *If $K = 1$, there are exactly two potentials $\boldsymbol{\pi}'$ and $\boldsymbol{\pi}''$ with $\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\boldsymbol{\pi}') = \mathbf{f}^{-1}(\boldsymbol{\pi}'')$ that are contained in nonambiguous, neighboring regions $R_{\mathbf{t}'}$ and $R_{\mathbf{t}''}$.*

iii. *If $K = 1$ and $\boldsymbol{\pi} \in R_{\mathbf{t}'}$, where $R_{\mathbf{t}'}$ is a nonambiguous, neighboring region of $R_{\mathbf{t}}$, there is another unique nonambiguous neighboring region $R_{\mathbf{t}''}$, and a potential $\boldsymbol{\pi}''$ with $\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\boldsymbol{\pi}'')$. In this case, $\boldsymbol{\pi}''$ and $R_{\mathbf{t}''}$ can be found in $\mathcal{O}(n^2)$ time. Given the inverse Laplacian matrix $\mathbf{L}^*_{\mathbf{t}'}$, the inverse Laplacian matrix $\mathbf{L}^*_{\mathbf{t}''}$ can be computed in $\mathcal{O}(n^2)$ time.*

Proof. Lemma 3 implies that there are $K + 1$ actively connected components in region $R_{\mathbf{t}}$. For every actively connected component $U \subseteq V$, denote by $\mathbf{1}_U$ be the indicator vector of the component, that is, $\mathbf{1}_U := (\mathbb{1}_{v \in U})_{v \in V} \in \mathbb{R}^n$. Then, $\mathbf{1}_U \in \ker(\mathbf{L}_{\mathbf{t}})$. We assume without loss of generality that $s \notin U$ (if $s \in U$, consider the component $V \setminus U$) for all components U we consider throughout this proof. This assumption ensures that $\mathbf{1}_U \in \Pi$. Given a potential $\boldsymbol{\pi} \in R_{\mathbf{t}}$ in an ambiguous region $R_{\mathbf{t}}$, we perform the following procedure:

1. Identify an arbitrary actively connected component $U \subset V$ (with $s \notin U$). Let $\mathbf{1}_U$ be the indicator vector of this component.

2. Compute $\mu \in \mathbb{R}$ with maximal absolute value $|\mu|$ such that the potential $\boldsymbol{\pi} + \mu \mathbf{1}_U$ is contained in the region $R_{\mathbf{t}}$ by solving the inequalities

$$\sigma_{e,t_e} \leq \boldsymbol{\gamma}_e^\top \boldsymbol{\pi} + \mu \boldsymbol{\gamma}_e^\top \mathbf{1}_U \leq \sigma_{e,t_e+1} \tag{17}$$

for all edges $e \in \delta(V')$ in the cut induced by the component V'. (By the definition of the actively connected component, all these edges are not actively connected.)

3. The potential $\boldsymbol{\pi} + \mu \mathbf{1}_U$ lies on the boundary to the neighboring region $R_{\mathbf{t}'}$ with the boundary edge e that induced the maximal μ in the previous computation. In this region, the boundary edge must now be actively connected, merging the actively connected component U with another actively connected component.

4. Now starting from $\boldsymbol{\pi} + \mu \mathbf{1}_U$ in region $R_{\mathbf{t}'}$, repeat Steps 2 and 3 until there is only one actively connected component using the merged actively connected component.

We claim that this procedure yields a potential $\boldsymbol{\pi}$ in a nonambiguous region $R_{\mathbf{t}'}$. By definition of $\mathbf{1}_U$, $\boldsymbol{\gamma}_e^\top \mathbf{1}_U = \pm \mathbf{1}$ if $e \in \delta(U)$ and $\boldsymbol{\gamma}_e^\top \mathbf{1}_U = 0$ otherwise. This implies two things. First, adding $\mathbf{1}_U$ to the potential $\boldsymbol{\pi}$ only changes the potential differences along edges in the cut $\delta(U)$. Because these edges are all inactive, that is, $c_{e,t_e} = 0$ for these edges, the induced flow does not change on any edge. Second, because the potential difference $\boldsymbol{\gamma}_e^\top \mathbf{1}_U$ is nonzero for all edges in the cut $\delta(U)$, there must be a maximal μ such that Inequalities (17) are satisfied.

Then, by the definition of the maximal μ, there must be at least one edge e^* such that $\boldsymbol{\gamma}_e^\top \boldsymbol{\pi} + \mu \boldsymbol{\gamma}_{e^*}^\top \mathbf{1}_U \in \{\sigma_{e^*,t}, \sigma_{e^*,t+1}\}$; that is, the potential $\boldsymbol{\gamma}_e^\top \boldsymbol{\pi} + \mu \boldsymbol{\gamma}_{e^*}^\top \mathbf{1}_U$ lies on the boundary induced by the edge e^*. Thus, there is the neighboring region $R_{\mathbf{t}'}$ with $\boldsymbol{\gamma}_e^\top \boldsymbol{\pi} + \mu \boldsymbol{\gamma}_{e^*}^\top \mathbf{1}_U \in R_{\mathbf{t}'}$. Because this region is e^*-neighboring to $R_{\mathbf{t}}$, the edge e^* must satisfy $c_{e^*,t'_{e^*}} \neq 0$. Hence, in $R_{\mathbf{t}'}$ the actively connected component U is now connected to another actively connected component U'. In particular, the number of connected components decreases by one, and thus, the rank of the Laplacian matrix $\mathbf{L}_{\mathbf{t}'}$ must be greater than the rank of $\mathbf{L}_{\mathbf{t}}$. Hence, after at most K iterations, this procedure yields a potential with the same flow in a nonambiguous region. If the same, growing connected component is used in all

iterations, Inequalities (17) need only be solved once for every edge, yielding an overall complexity of $\mathcal{O}(n^2)$. This proves (i).

For (ii), we observe that if $K = 1$, there are exactly two actively connected components, and only one of them containing the vertex s. Thus, there is only one direction $\mathbf{1}_U$ along which the flow does not change. Using either $\mathbf{1}_U$ or $-\mathbf{1}_U$ in (17) yields two potentials in neighboring regions inducing the same flow. These regions must have fewer actively connected components and hence be nonambiguous. Furthermore, there cannot be more such potentials as their existence would imply the existence of more actively connected components in the region $R_{\mathbf{t}}$.

Finally, if we are already given one of these potentials as described in (iii), there is exactly one other unique potential $\boldsymbol{\pi}''$ in a nonambiguous region. This can be found by solving (17) in $\mathcal{O}(n^2)$. The new Laplacian can be computed by using the rank 2 update

$$\mathbf{L}_{\mathbf{t}''} = \mathbf{L}_{\mathbf{t}'} - c_{e_1,t'_{e_1}} \boldsymbol{\gamma}_{e_1} \boldsymbol{\gamma}_{e_1}^\top + c_{e_2,t''_{e_2}} \boldsymbol{\gamma}_{e_2} \boldsymbol{\gamma}_{e_2}^\top,$$

where the edges e_1 and e_2 are the boundary edges between the regions $R_{\mathbf{t}}$ and $R_{\mathbf{t}'}$, as well as $R_{\mathbf{t}}$ and $R_{\mathbf{t}''}$. The inverse $L^*_{\mathbf{t}''}$ can be updated using the Sherman-Morrison-Woodbury (Harville [27, theorem 18.2.8]) formula in $\mathcal{O}(n^2)$ time. □

The main message of Theorem 4 is that ambiguous regions can be skipped by jumping to a different potential in a neighboring, nonambiguous region without changing the flow. For the parametric algorithm to be well defined, we need two additional properties that we prove in the following lemma:

i. When performing a pivoting step, the next region should be not "too ambiguous," that is, should be a region with $K = n - 1 - \operatorname{rank}(\mathbf{L}_{\mathbf{t}}) = 1$. In this case, Theorem 4(iii) ensures that there is a unique nonambiguous region to jump to.

ii. The region before and the region after the ambiguous region that is skipped should have compliant potential directions $\Delta\pi_{\mathbf{t}}$ in the sense that the potential direction in the region before the ambiguous region should point toward and the potential direction in the region after the ambiguous region should point away from the ambiguous region. This prevents the algorithm from jumping over the same ambiguous region over and over again.

Lemma 8. *Let $R_{\mathbf{t}}$ be a feasible, nonambiguous region and let $R_{\mathbf{t}'}$ be the next region after a pivoting step as defined in Corollary 2. Then,*

i. *We have* $\operatorname{rank}(\mathbf{L}_{\mathbf{t}'}) \geq n - 2$.

ii. *If $R_{\mathbf{t}'}$ is ambiguous, there is a unique nonambiguous region $R_{\mathbf{t}''}$ such that the potentials $\boldsymbol{\pi} := \boldsymbol{\pi}_{\mathbf{t}} + \lambda_{\mathbf{t}}^{\max} \Delta\boldsymbol{\pi}_{\mathbf{t}}$ and $\boldsymbol{\pi}'' := \pi_{\mathbf{t}''} + \lambda_{\mathbf{t}''}^{\min} \Delta\pi_{\mathbf{t}''}$ induce the same flow, that is, $\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\boldsymbol{\pi}'')$. Thus, in particular $\lambda_{\mathbf{t}}^{\max} = \lambda_{\mathbf{t}''}^{\min}$.*

Proof. Because the regions $R_{\mathbf{t}}$ and $R_{\mathbf{t}'}$ are neighboring, the Laplacian matrices satisfy (see also proof of Theorem 2) $\mathbf{L}_{\mathbf{t}'} = \mathbf{L}_{\mathbf{t}} + \Delta c_e \boldsymbol{\gamma}_e \boldsymbol{\gamma}_e^\top$, that is, the matrices can be obtained from each other by a rank 1 update. This implies, that $\operatorname{rank}(\mathbf{L}_{\mathbf{t}'}) \geq \operatorname{rank}(\mathbf{L}_{\mathbf{t}}) - 1$, proving (i).

By assumption and (i), the potential $\boldsymbol{\pi}$ and the region $R_{\mathbf{t}'}$ satisfy the conditions of Theorem 4(iii) implying the existence of a unique nonambiguous region $R_{\mathbf{t}''}$ and a potential $\boldsymbol{\pi}'' \in R_{\mathbf{t}''}$ with $\mathbf{f}^{-1}(\boldsymbol{\pi}) = \mathbf{f}^{-1}(\boldsymbol{\pi}'')$. It remains to be shown that $\boldsymbol{\pi}'' := \pi_{\mathbf{t}''} + \lambda_{\mathbf{t}''}^{\min} \Delta\pi_{\mathbf{t}''}$. To this end, let e_1^* be the boundary edge between region $R_{\mathbf{t}}$ and $R_{\mathbf{t}'}$ (i.e., the edge inducing $\lambda_{\mathbf{t}}^{\max}$). Let e_2^* be the boundary edge between $R_{\mathbf{t}'}$ and $R_{\mathbf{t}''}$ (i.e., the edge inducing the μ with maximal absolute value satisfying (17) in the proof of Theorem 4). It may be that $e_1^* = e_2^*$. Furthermore, let $U \subset V$ be the unique actively connected component with $s \notin U$ in $R_{\mathbf{t}'}$ as defined in the proof of Theorem 4 and $\mathbf{1}_U$ be the indicator vector of the component U. Finally, let μ be the value as defined in the proof of Theorem 4.

All edges in the cut $\delta(U)$ are not active in $R_{\mathbf{t}'}$ as U. U is an actively connected component in $R_{\mathbf{t}'}$ by definition. Because e_1^* is the boundary edge between $R_{\mathbf{t}}$ and $R_{\mathbf{t}'}$, we infer that the edge e_1^* must be the only active edge in this cut in region $R_{\mathbf{t}}$ (the value c_e is unchanged for all other edges). The same is true for the edge e_2^* in the region $R_{\mathbf{t}''}$.

Claim 1. We have $\operatorname{sgn}(\boldsymbol{\gamma}_{e_1^*}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}}) = \operatorname{sgn}(\mathbf{1}_U^\top \mathbf{q}\ \boldsymbol{\gamma}_{e_1^*}^\top \mathbf{1}_U)$.

Proof of Claim 1. We observe that $\mathbf{1}_U^\top \boldsymbol{\gamma}_e = 0$ if $e \notin \delta(U)$. In the region $R_{\mathbf{t}}$, e_1^* is the only active edge (i.e., the only edge with $c_{e,t_e} \neq 0$) in the cut $\delta(U)$. Thus, we obtain

$$\mathbf{1}_U^\top \mathbf{q} = \mathbf{1}_U^\top \mathbf{L}_{\mathbf{t}} \Delta\boldsymbol{\pi}_{\mathbf{t}} = \mathbf{1}_U^\top \boldsymbol{\Gamma} \mathbf{C}_{\mathbf{t}} \boldsymbol{\Gamma}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}} = \sum_{e \in E} \mathbf{1}_U^\top \boldsymbol{\gamma}_e c_{e,t_e} \boldsymbol{\gamma}_e^\top \Delta\boldsymbol{\pi}_{\mathbf{t}} = \mathbf{1}_U^\top \boldsymbol{\gamma}_{e_1^*} c_{e_1^*,t_{e_1^*}} \boldsymbol{\gamma}_{e_1^*}^\top \Delta\boldsymbol{\pi}_{\mathbf{t}}.$$

With $c_{e_1^*,t_{e_1^*}} > 0$, the claim follows. △

Claim 2. We have $\operatorname{sgn}(\mu) = \operatorname{sgn}(\mathbf{1}_U^\top \mathbf{q})$.

Proof of Claim 2. Because we assume nondegeneracy, we have that $\mu \neq 0$. By the definition of $\lambda_{\mathfrak{t}}^{\max}$ of the region $R_{\mathfrak{t}}$ and the fact that $R_{\mathfrak{t}}$ and $R_{\mathfrak{t}'}$ are neighboring with boundary edge e_1^*, we get $\boldsymbol{\gamma}_{e_1^*}^\top \boldsymbol{\pi} = \sigma_{e_1^*, t'_{e_1^*}} = \sigma_{e_1^*, t_{e_1^*}+1}$ if and only if $\boldsymbol{\gamma}_{e_1^*}^\top \Delta\boldsymbol{\pi}_{\mathfrak{t}} > 0$ and $\boldsymbol{\gamma}_{e_1^*}^\top \boldsymbol{\pi} = \sigma_{e_1^*, t'_{e_1^*}+1} = \sigma_{e_1^*, t_{e_1^*}}$ if and only if $\boldsymbol{\gamma}_{e_1^*}^\top \Delta\boldsymbol{\pi}_{\mathfrak{t}} < 0$, where $\boldsymbol{\pi} := \boldsymbol{\pi}_{\mathfrak{t}} + \lambda_{\mathfrak{t}}^{\max} \Delta\boldsymbol{\pi}_{\mathfrak{t}}$ is the potential on the boundary between $R_{\mathfrak{t}}$ and $R_{\mathfrak{t}'}$. This implies that μ can be nonzero and satisfy Equation (17) only if the term $\mu \boldsymbol{\gamma}_{e_1^*}^\top \mathbf{1}_U$ and $\boldsymbol{\gamma}_{e_1^*}^\top \Delta\boldsymbol{\pi}_{\mathfrak{t}}$ have the same sign. Together with Claim 1, $\operatorname{sgn}(\mu) = \operatorname{sgn}(\mathbf{1}_U^\top \mathbf{q})$ follows. △

Claim 3. We have $\operatorname{sgn}(\mu\, \boldsymbol{\gamma}_{e_2^*}^\top \mathbf{1}_U) = \operatorname{sgn}(\boldsymbol{\gamma}_{e_2^*}^\top \Delta\boldsymbol{\pi}_{\mathfrak{t}''})$.

Proof of Claim 3. With Claim 2, we get $\operatorname{sgn}(\mu\, \boldsymbol{\gamma}_{e_2^*}^\top \mathbf{1}_U) = \operatorname{sgn}(1_U^\top \mathbf{q}\, \boldsymbol{\gamma}_{e_2^*}^\top \mathbf{1}_U)$. Using the same argument as in the proof of Claim 1 for the region $R_{\mathfrak{t}''}$ and the boundary edge e_2^*, we get $\operatorname{sgn}(\boldsymbol{\gamma}_{e_2^*}^\top \Delta\pi_{\mathfrak{t}''}) = \operatorname{sgn}(\mathbf{1}_U^\top \mathbf{q}\, \boldsymbol{\gamma}_{e_2^*}^\top \mathbf{1}_U)$. △

By the definition of μ, we have that $\boldsymbol{\pi}'' = \boldsymbol{\pi} + \mu \mathbf{1}_U$. Because $\boldsymbol{\pi}''$ is the boundary potential of the regions $R_{\mathfrak{t}'}$ and $R_{\mathfrak{t}''}$, either $\boldsymbol{\gamma}_{e_2^*}^\top \boldsymbol{\pi}'' = \sigma_{e_2^*, t'_{e_2^*}}$ or $\boldsymbol{\gamma}_{e_2^*}^\top \boldsymbol{\pi}'' = \sigma_{e_2^*, t'_{e_2^*}+1}$. The former is true if and only if $\mu < 0$, the latter is true if and only if $\mu > 0$. With Claim 3, we obtain $\boldsymbol{\gamma}_{e_2^*}^\top \boldsymbol{\pi}'' = \sigma_{e_2^*, t''_{e_2^*}+1}$ if and only if $\boldsymbol{\gamma}_{e_2^*}^\top \Delta\boldsymbol{\pi}_{\mathfrak{t}''} > 0$ and $\boldsymbol{\gamma}_{e_2^*}^\top \boldsymbol{\pi}'' = \sigma_{e_2^*, t''_{e_2^*}}$ if and only if $\boldsymbol{\gamma}_{e_2^*}^\top \Delta\pi_{\mathfrak{t}''} < 0$. The definition of the values $\overline{\lambda}$ and $\underline{\lambda}$, as well as the definition of $\lambda_{\mathfrak{t}''}^{\min}$ in Theorem 1 yield that the respective breakpoint contributes to the maximum in the computation of $\lambda_{\mathfrak{t}''}^{\min}$. Thus, $\boldsymbol{\pi}'' = \boldsymbol{\pi}_{\mathfrak{t}''} + \lambda_{\mathfrak{t}''}^{\min} \Delta\pi_{\mathfrak{t}''}$. Because the potentials $\boldsymbol{\pi}$ and $\boldsymbol{\pi}''$ induce the same flow, we in particular have $\lambda_{\mathfrak{t}}^{\max} = \lambda_{\mathfrak{t}''}^{\min}$. □

Example 3. Consider the graph from Figure 7(a) with the discontinuous marginal cost functions on the edges. As an example, the marginal cost function $f_3(x) = x$ for $x \leq 3$ and $f_3(x) = 3x - 4$ for $x > 3$ for edge e_3 has the inverse function $f_3^{-1}(x) = x$ if $x \leq 3$, $f_3^{-1}(x) = 3$ if $3 < x \leq 5$ and $f_3^{-1}(x) = \frac{x+4}{3}$ if $x > 5$. Thus, all potential differences $3 < \pi_t - \pi_s < 5$ induce the same flow $x_{e_3} = 3$. In all regions where $3 < \pi_t - \pi_s < 5$, the edge e_3 is inactive. Figure 7(b) shows the potential space for the given graph with $\pi_s = 0$ fixed. For all potentials $3 < \pi_t - \pi_s < 5$ (i.e., all potentials in between the vertical dashed lines), edge e_3 is inactive. Likewise, for all potentials in between the horizontal dashed lines, e_1 is inactive. These lines mark the potentials where $\boldsymbol{\gamma}_{e_1}^\top \boldsymbol{\pi} = \pi_v - \pi_s = f_1^-(2) = 2$ and $\boldsymbol{\gamma}_{e_1}^\top \boldsymbol{\pi} = \pi_v - \pi_s = f_1^+(2) = 3$. For all potentials in between the dashed diagonal lines, the edge e_2 is inactive. The regions where these areas intersect are ambiguous regions. For example, for all potentials in the lower gray region, the edges e_2 and e_3 are inactive. In this region, there are two actively connected components $\{s, v\}$ and $\{t\}$. Increasing or decreasing the potential π_t does not change the flows on the incident edges e_2, e_3. This direction is indicated by the zigzag arrow. In the middle region marked in dark gray, all edges are inactive; thus, there are three actively connected components $\{s\}$, $\{v\}$, and $\{t\}$. The potentials π_v and π_t (π_s is fixed to 0) can be changed

Figure 7. Network with discontinuous marginal cost functions from Example 3. (a) The graph with the discontinuous marginal cost functions. (b) The potential space for $\pi_s = 0$. The shaded regions are ambiguous, the zigzig arrows indicate the directions along which the induced flow does not change in these regions. The thick line represents the solution curve for demands $(q_s, q_v, q_t) = \lambda(-1, 0, 1)$.

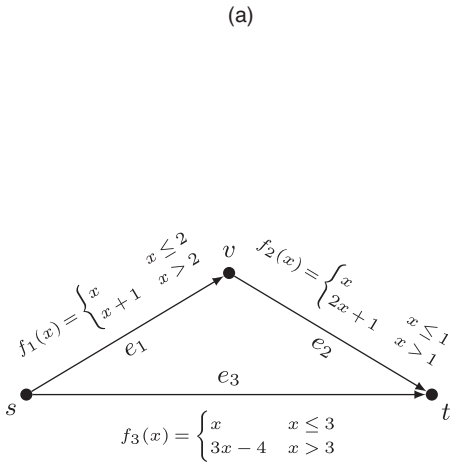

Graph with marginal cost functions.

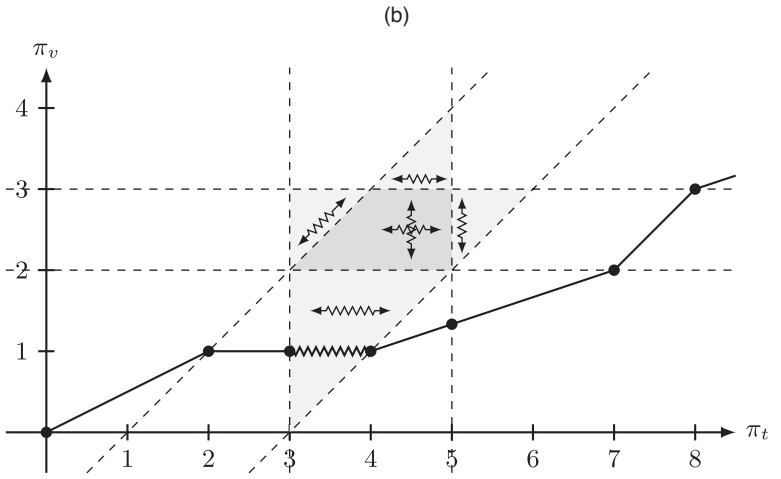

Potential space for $\pi_s = 0$ with solution curve.

arbitrarily without changing the induced flow. Hence, in this region, two directions along which the flow does not change exist.

Finally, consider the demand vector $\mathbf{q} = (q_s, q_v, q_t) = (-1, 0, 1)$, that is, the demand vector of the s-t flow for demand 1. The thick line in Figure 7(a) represents the solution curve of the parametric solution for demands $\lambda\mathbf{q}, \lambda \geq 0$, that is, for all s-t flows with demand λ. The solution curve starts in the potential $\boldsymbol{\pi} = (\pi_s, \pi_v, \pi_t) = (0, 0, 0)$. In the potential $\boldsymbol{\pi} = (0, 1, 2)$, the induced flow on edge e_2 is $x_{e_2} = f_2^{-1}(2-1) = 1$ and reaches the breakpoint of the marginal cost function. The edge becomes inactive, and all further flow takes the lower edge e_3. Therefore, the potential π_v is fixed while the potential π_t grows (reflected in the horizontal part of the solution curve), until the edge e_3 becomes inactive as well. The solution curve enters the ambiguous region and jumps along the zigzag line to the end of the ambiguous region. The induced flow of the potentials $\boldsymbol{\pi} = (0, 1, 3)$ and $\boldsymbol{\pi} = (0, 1, 4)$ is indeed the same flow $\mathbf{x} = (1, 1, 3)$. At the latter potential, e_2 becomes active again, and the solution curve proceeds normally.

4.2. Modeling Directed Edges and Edge Capacities

One major advantage of allowing discontinuous marginal cost functions is that this allows modeling of edge capacities, that is, hard constraints on the edge flow of the form $l_e(x_e) \leq x_e \leq u_e(x_e)$. In particular, we can model directed edges via lower edge flow constraints of the form $x_e \geq 0$, restricting the flow direction to the edge direction. In order to model edge capacities, we need to admit slightly more general marginal cost functions as we assume that the marginal cost functions are functions $f_e : \mathbb{R} \to \mathbb{R} \cup \{-\infty, +\infty\}$ that are strictly increasing, piecewise linear and satisfy $f_e^-(0) \leq 0 \leq f_e^+(0)$; that is, in contrast to the last section, we also admit that the marginal costs can have infinite values. (With some abuse of terminology, we consider such a function to be strictly increasing if it is strictly increasing on the maximal interval where it takes finite values.) The main claim of this section is that $f_e(x) = -\infty$ for all $x < l_e$ models a lower capacity of l_e and $f_e(x) = +\infty$ for all $x > u_e$ models an upper capacity of u_e[2] (Figure 8).

Theorem 5. *Let $\mathbf{q} \in \mathbb{R}^n$ be a demand vector, and for all $e \in E$ let $l_e \leq 0 \leq u_e$ be upper and lower edge capacities such that there is a feasible flow $\mathbf{x}$ with $\boldsymbol{\Gamma}\mathbf{x} = \mathbf{q}$ and $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$. Then, for strictly increasing marginal cost function $f_e : [l_e, u_e] \to \mathbb{R}$ a flow $\mathbf{x}$ is the unique optimal solution to*

$$\min \quad \sum_{e \in E} F_e(x_e) \quad \text{s.t.} \quad \boldsymbol{\Gamma}\mathbf{x}=\mathbf{q} \quad \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}. \tag{18}$$

if and only if $\mathbf{x}$ is the unique optimal solution to

$$\min \quad \sum_{e \in E} \tilde{f}_e(x_e) \quad \text{s.t.} \quad \boldsymbol{\Gamma}\mathbf{x}=\mathbf{q}, \tag{19}$$

where $\tilde{F}(x) := \int_0^x \tilde{f}_e(s)\,\mathrm{d}s$ with

$$\tilde{f}_e(x) = \begin{cases} -\infty & \text{if } x < l_e, \\ f_e(x) & \text{if } l_e \leq x \leq u_e, \\ +\infty & \text{if } x > u_e. \end{cases}$$

Figure 8. A discontinuous characteristic function f_e modeling a lower edge capacity of l_e and an upper edge capacity of u_e. (a) The characteristic function with $f_e(x) = -\infty$ for $x < l_e$ and $f_e(x) = \infty$ for $x > u_e$. (b) The (pseudo-)inverse characteristic function f_e^{-1}.

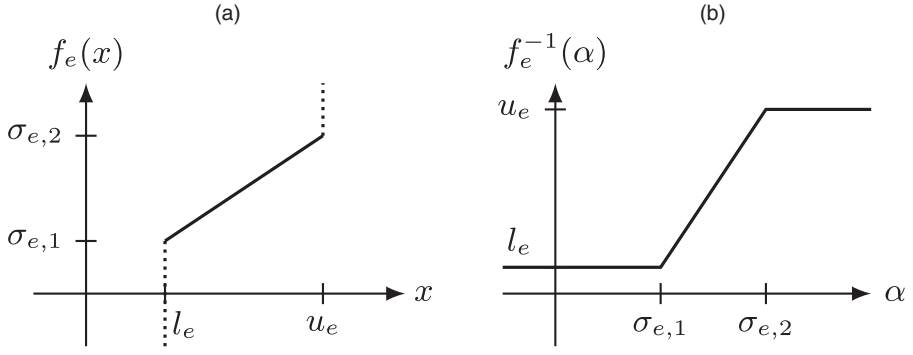

Characteristic function. Inverse characteristic function.

Proof. Every $\mathbf{x}$ that is feasible for (18) is also feasible for (19) and has a finite objective value in both systems. Because every $\mathbf{x}$ that is feasible in (19) but not feasible in (18) has a nonfinite objective value in (19), the claim follows. □

The marginal cost functions $\tilde{f}_e$ from Theorem 5 can be treated like the discontinuous marginal cost from the previous section. The statements from Theorem 4 are still valid with a minor restriction: For every potential in an ambiguous region, we still can find another potential in a nonambiguous region inducing the same flow. However, we are not guaranteed to always find at least two such potentials. This means it may happen that the algorithm, when encountering an ambiguous region, cannot jump to another potential and proceed as before—however, as it turns out, in this case the algorithm has found a maximal flow. We summarize these results in the following lemma.

Lemma 9. *Assume the characteristic functions may also take values $+\infty$ or $-\infty$. As in Theorem 4, let $\boldsymbol{\pi} \in R_{\mathbf{t}}$ be a potential in an ambiguous region $R_{\mathbf{t}}$. Denote by $K := n - 1 - \operatorname{rank}(\mathbf{L}_{\mathbf{t}})$ the dimension of the kernel of $\mathbf{L}_{\mathbf{t}}$ in Π. Then,*

i. *The statement of Theorem 4(i) holds.*

ii. *If $K = 1$, either the statement of Theorem 4(iii) holds or there is a vertex cut $V' \subset V$ such that $x_e = u_e$ for all outgoing edges $e \in \delta^+(V')$ and $x_e = l_e$ for all incoming edges $e \in \delta^-(V')$, where $\mathbf{x} = \mathbf{C}_{\mathbf{t}}\boldsymbol{\Gamma}^\top\boldsymbol{\pi} - \mathbf{d}_{\mathbf{t}}$ is the flow induced by $\boldsymbol{\pi}$.*

Proof. The only difference to the situation in Theorem 4 is that the for every cut edges $e \in \delta(U)$, one of the values σ_{e,t_e} or σ_{e,t_e+1} in (17) may be infinite. However, statement (i) of Theorem 4 still holds, because even if Equation (17) has no maximal solution μ, there must be a maximal μ if $\mathbf{1}_U$ is replaced by $-\mathbf{1}_U$. Because at most, one of the values σ_{e,t_e} or σ_{e,t_e+1} can be infinite in one region for every edge, there must exists a finite μ for either $+1_U$ or -1_U.

However, it may happen that statement (iii) of Theorem 4 does not hold anymore, because by the same arguments as before, only the existence of one potential in one nonambiguous region is guaranteed. However, this can only be the case if for all cut edges $e \in \delta(U)$, the respective values σ_{e,t_e} or σ_{e,t_e+1} are infinite. However, this means by definition that all edges in the cut must either be at maximum or minimum capacity, proving the claim. □

4.3. Nonhomogeneous Marginal Costs

Thus far, we made the assumption that all marginal cost functions are homogeneous, that is, $f_e^-(0) \le 0 \le f_e^+(0)$. This assumption ensures that the potential $\boldsymbol{\pi} = \mathbf{0}$ is the optimal potential of the zero flow $\mathbf{x} = \mathbf{0}$. Thus, we can start the algorithm with $\lambda = 0$. In case of discontinuous marginal cost functions, the starting region may be ambiguous, but Theorem 4(i) enables us to find a potential in a neighboring nonambiguous region that induces the zero flow as well in $\mathcal{O}(n^2)$ time. In fact, it is not hard to show that an initial potential in a nonambiguous region can be found by a simple shortest path computation, for example, with Dijktstra's algorithm, in $\mathcal{O}(n^2)$ time. However, dropping the homogeneity assumption, the zero potential does not necessarily induce the zero flow. Moreover, it may be the case that the zero flow is not even the optimal solution for $\lambda = 0$ or that there is no feasible flow for $\lambda = 0$ at all.

The homogeneity of the marginal cost functions ensures that the graph of the marginal cost function f_e of every edge e lies only in the first and the third quadrant. From the perspective of electrical networks, electrical components with this property are called *passive components*; they consume power from the circuit, for example, resistors and motors. On the other hand, electrical components with characteristics whose graph lies in the second and fourth quadrant are called *active components*; these components actually produce power, for example, batteries and generators.

Our strategy to circumvent the lack of a starting potential and starting flow for our algorithm is to add a first phase where a feasible flow for minimal λ is found (if it exists), and a corresponding potential is computed. To this end, we use the following steps.

Finding $\lambda^{\min}$: Find the minimal $\lambda^{\min} \ge 0$ such that there is a feasible flow $\mathbf{x}$ for demands $\lambda^{\min}\mathbf{q}$. This step is necessary only in the presence of hard edge capacities $\mathbf{l} \le \mathbf{x} \le \mathbf{u}$, as otherwise $\lambda^{\min} = 0$.

Phase I: Start with the zero potential $\boldsymbol{\pi}^0 = \mathbf{0}$ in some initial region $R_{\mathbf{t}}$. If $R_{\mathbf{t}}$ is ambiguous, find a potential and a region inducing the same flow contained in a nonambiguous region, as described in Theorem 4(i). Then use the algorithm to compute a potential $\boldsymbol{\pi}^{\min}$ in a region $R_{\mathbf{t}^{\min}}$ that induces a flow $\mathbf{x} = \mathbf{C}_{\mathbf{t}^{\min}}\boldsymbol{\Gamma}^\top\boldsymbol{\pi} - \mathbf{d}_{\mathbf{t}^{\min}}$ for minimal feasible demands $\lambda^{\min}\mathbf{q}$.

Phase II: Start the parametric computation with the initial potential $\boldsymbol{\pi}^{\min}$.

We proceed by describing these steps in more detail.

Finding λ^{min}. In the case of hard edge capacities of the form $\mathbf{l} \le \mathbf{x} \le \mathbf{u}$, the problem may not be feasible for all demands $\lambda\mathbf{q}, \lambda \ge 0$, but only for an maximal interval $\Lambda = [\lambda^{\min}, \lambda^{\max}]$ for some $\lambda^{\min} \in \mathbb{R}_{\ge 0}$, $\lambda^{\max} \in \mathbb{R}_{\ge 0} \cup \{\infty\}$. The convex combination of feasible flows is feasible implying that Λ is convex as claimed. To proceed, we need to

compute $\lambda^{\min}$ before starting the parametric computation. The minimal λ can be found by solving the linear program

$$\min \quad \lambda \quad \text{s.t.} \quad \mathbf{\Gamma x}=\lambda \mathbf{q}, \quad \mathbf{l} \le \mathbf{x} \le \mathbf{u}, \quad \lambda \ge 0.$$

This linear program is either infeasible—in this case, there is no $\lambda \ge 0$ such that any feasible flow for demand $\lambda \mathbf{q}$ exists—or the linear program yields the minimal λ for which a feasible flow exists.

Phase I. For nonhomogeneous marginal cost functions, the zero potential $\boldsymbol{\pi} = \mathbf{0}$ need not induce a flow for the minimal feasible demand $\lambda^{\min}\mathbf{q}$. In order to find an initial solution, one could solve the nonparametric problem for fixed demand $\lambda^{\min}\mathbf{q}$ with an algorithm for nonparametric mincost flows. Here, we describe how our parametric algorithm can be used to find an initial solution in a first phase. First, start with the zero potential $\boldsymbol{\pi}^0 = \mathbf{0} \in R_{\mathbf{f}^0}$ (if $R_{\mathbf{f}^0}$ is ambiguous, use another potential in a nonambiguous region, e.g., the potential described in Theorem 4(i)). This potential induces some flow $\mathbf{x}^0 = \mathbf{C}_{\mathbf{f}^0}\mathbf{\Gamma}^\top\boldsymbol{\pi}^0 - \mathbf{d}_{\mathbf{f}^0}$ that is a feasible for some demand $\mathbf{q}^0 = \mathbf{\Gamma x}$ different from the minimal demand $\lambda^{\min}\mathbf{q}$. Let $\Delta\mathbf{q} := \lambda^{\min}\mathbf{q} - \mathbf{q}^0$ be the difference of these demands. We start the parametric algorithm with demand $\Delta\mathbf{q}$ starting from $\boldsymbol{\pi}^0$ and obtain a solution curve $\lambda \mapsto \tilde{\boldsymbol{\pi}}(\lambda)$ where $\tilde{\boldsymbol{\pi}}(\lambda)$ is the potential for the electrical flow for demand $\mathbf{q}^0 + \lambda\Delta\mathbf{q}$. Hence, the potential $\tilde{\boldsymbol{\pi}}(1)$ is a potential of the electrical flow for demand $\lambda^{\min}\mathbf{q}$ and thus the initial potential for the actual parametric computation.

Phase II. Equipped with the starting flow $\lambda^{\min}$ and the starting potential $\tilde{\boldsymbol{\pi}}$ from Phase I, we use the basic parametric algorithm to compute the parametric solution.

We illustrate the different phases of the algorithm with the following example.

Example 4. Consider the graph from Figure 9(a) with the marginal cost functions given next to the edges. We compute a parametric solution for demands $\lambda\mathbf{q}$ with $\mathbf{q} = (q_s, q_v, q_t) = (-1, 0, 1)$, that is, s-t flows for flow rates λ. Because we have no hard edge capacities, the minimal feasible demand is $\lambda^{\min}\mathbf{q} = \mathbf{0}$ because $\lambda^{\min} = 0$. However, for the characteristic function f_1 of the edge e_1, we have $f_1(0) = 4 \neq 0$. Thus, the zero potential $\boldsymbol{\pi}^0 = \mathbf{0}$ does not induce a flow for demand $\lambda^{\min} = 0$. Rather, the zero potential induces the flow $\mathbf{x}^0 = (-6, 0, 0)$ (see first graph in Figure 9(c)). This

Figure 9. Input and Phase I and II solution of nonhomogeneous electrical flow problem from Example 4.

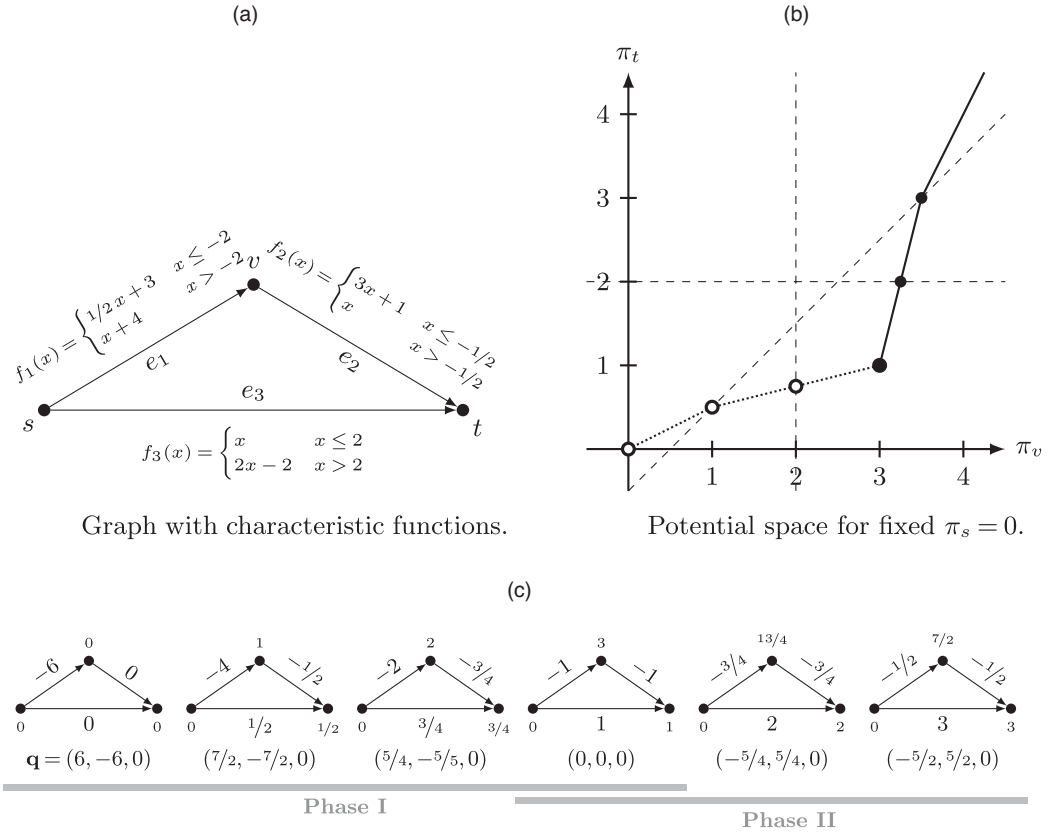

Graph with characteristic functions.

Potential space for fixed $\pi_s = 0$.

Table 2. Iterations for the parametric solution from Example 5. For the lexicographic rule used in iteration $i = 0$, the first component is ignored for the lexicographic comparison.

i	$\mathbf{t}^i$	$\mathbf{L}_{\mathbf{t}^i}$	$\mathbf{L}^*_{\mathbf{t}^i}$	$\pi_{\mathbf{t}^i}$	$\Delta\pi_{\mathbf{t}^i}$	$\lambda^{\min}_{\mathbf{t}^i}$	$\lambda^{\max_{\mathbf{t}^i}}$	$E^*_{\mathbf{t}^i}$	Lex. rule
0	$\begin{bmatrix}2\\1\\2\\1\\2\end{bmatrix}$	$\begin{bmatrix}1/2 & -1/2 & 0 & 0\\ -1/2 & 3/2 & -1 & 0\\ 0 & -1 & 3/2 & -1/2\\ 0 & 0 & -1/2 & 1/2\end{bmatrix}$	$\begin{bmatrix}0&0&0&0\\0&2&2&2\\0&2&3&3\\0&2&3&5\end{bmatrix}$	$\begin{bmatrix}0\\0\\0\\0\end{bmatrix}$	$\begin{bmatrix}0\\2\\3\\5\end{bmatrix}$	0	1	$\{e_2, e_4\}$	$\begin{bmatrix}*\\0\\-1/3\\0\end{bmatrix}_{\mathbf{m}_{e_2,\mathbf{t}^0}} \prec_{\text{LEX}} \begin{bmatrix}*\\1/3\\0\\-1/3\end{bmatrix}_{\mathbf{m}_{e_4,\mathbf{t}^0}}$
1	$\begin{bmatrix}2\\2\\2\\1\\2\end{bmatrix}$	$\begin{bmatrix}3/2 & -1/2 & -1 & 0\\ -1/2 & 3/2 & -1 & 0\\ -1 & -1 & 5/2 & -1/2\\ 0 & 0 & -1/2 & 1/2\end{bmatrix}$	$\begin{bmatrix}0&0&0&0\\0&1&1/2&1/2\\0&1/2&3/4&3/4\\0&1/2&3/4&9/4\end{bmatrix}$	$\begin{bmatrix}0\\3/2\\5/5\\11/4\end{bmatrix}$	$\begin{bmatrix}0\\1/2\\3/4\\9/4\end{bmatrix}$	1	1	$\{e_4\}$	–
2	$\begin{bmatrix}2\\2\\2\\2\\2\end{bmatrix}$	$\begin{bmatrix}3/2 & -1/2 & -1 & 0\\ -1/2 & 5/2 & -1 & -1\\ -1 & -1 & 5/2 & -1/2\\ 0 & -1 & -1/2 & 3/2\end{bmatrix}$	$\begin{bmatrix}0&0&0&0\\0&14/15&8/15&4/5\\0&8/15&11/15&3/5\\0&4/5&3/5&7/5\end{bmatrix}$	$\begin{bmatrix}0\\6/5\\12/5\\18/5\end{bmatrix}$	$\begin{bmatrix}0\\4/5\\3/5\\7/5\end{bmatrix}$	1	6	$\{e_3\}$	–
3	$\begin{bmatrix}2\\2\\1\\2\\2\end{bmatrix}$	$\begin{bmatrix}3/2 & -1/2 & -1 & 0\\ -1/2 & 3/2 & 0 & -1\\ -1 & 0 & 3/2 & -1/2\\ 0 & -1 & -1/2 & 3/2\end{bmatrix}$	$\begin{bmatrix}0&0&0&0\\0&4/3&1/3&1\\0&1/3&5/6&1/2\\0&1&1/2&3/2\end{bmatrix}$	$\begin{bmatrix}0\\0\\3\\3\end{bmatrix}$	$\begin{bmatrix}0\\1\\1/2\\3/2\end{bmatrix}$	6	∞	$\emptyset$	–

flow satisfies the demands $\mathbf{q}^0 = (6, -6, 0)$. Thus, we compute the Phase I solution curve $\tilde{\boldsymbol{\pi}}$ starting in the zero potential for the demand $\Delta\mathbf{q} := 0 - \mathbf{q}^0 = (-6, 6, 0)$. The Phase I solution curve is depicted by the dotted curve in Figure 9(b). For $\lambda = 1$, the solution curve yields the potential $\tilde{\boldsymbol{\pi}}(1) = (0, 3, 1)$, inducing the flow $\mathbf{x} = (-1, -1, 1)$, that is, the electrical flow for zero demand (see the fourth graph in Figure 9(c)). This potential is the initial potential for the actual parametric computation in Phase II. The Phase II solution curve is depicted as the solid curve in Figure 9(b).

We illustrate the interplay of Phases I and II with the following example.

Example 5. Consider the mincost flow problem from Example 1 with the graph and cost functions depicted in Figure 1. In order to solve the parametric mincost problem, we transform the problem into a parametric electrical flow problem as previously described. The characteristic functions are $f_{e_1}(x) = f_{e_5}(x) = 2x$, $f_{e_2}(x) = f_{e_4}(x) = x + 3$, and $f_{e_3}(x) = x$ for $x \geq 0$ and $f_e = -\infty$ for $x < 0$ and $e \in E$. Therefore, all characteristics have one breakpoint at $x = 0$. The initial potential $\boldsymbol{\pi} = \mathbf{0}$ lies in the nonambiguous region $R_{\mathbf{t}^0}$ with $\mathbf{t}^0 = (2,1,2,1,2)^\top$. (Note that $\boldsymbol{\pi} = \mathbf{0}$ is also contained in other regions, e.g., $R_{(1,1,1,1,1)^\top}$, but all these other regions are ambiguous, and Theorem 4 would yield region $R_{\mathbf{t}^0}$ when starting in one of the other regions.) Starting from this region, Table 2 lists all relevant quantities for the parametric solution. In the first iteration $i = 0$, there are two boundary edges requiring the lexicographic rule. Recall that for the lexicographic comparison, the first components of the vectors $\mathbf{m}_{e,\mathbf{t}}$ are ignored, indicated by the star. The second iteration $i = 1$ produces a line segment of zero length because $\lambda^{\min}_{\mathbf{t}^1} = \lambda^{\max}_{\mathbf{t}^1} = 1$. Hence, this region does not contribute to the solution. Overall, we obtain the solution curve

$$\boldsymbol{\pi}(\lambda) = \begin{cases} \lambda\,(0,2,3,5)^\top & 0 \leq \lambda < 1, \\ (0, \frac{6}{5}, \frac{12}{5}, \frac{18}{5})^\top + \lambda\,(0, \frac{4}{5}, \frac{3}{5}, \frac{7}{5})^\top & 1 \leq \lambda < 6, \\ (0,0,3,3)^\top + \lambda\,(0,1, \frac{1}{2}, \frac{3}{2})^\top & \lambda \geq 6. \end{cases}$$

Applying the function $\mathbf{f}^{-1}$, we obtain the solution

$$\mathbf{x}(\lambda) = \mathbf{f}^{-1}(\boldsymbol{\pi}(\lambda)) = \begin{cases} \lambda\,(1,0,1,0,1)^\top & 0 \leq \lambda < 1, \\ (\frac{3}{5}, -\frac{3}{5}, \frac{6}{5}, -\frac{3}{5}, \frac{3}{5})^\top + \lambda\,(\frac{2}{5}, \frac{3}{5}, -\frac{1}{5}, \frac{3}{5}, \frac{2}{5})^\top & 1 \leq \lambda < 6, , \\ \lambda\,(\frac{1}{2}, \frac{1}{2}, 0, \frac{1}{2}, \frac{1}{2})^\top & \lambda \geq 6 \end{cases}$$

which is also depicted in Figure 1(b).

5. Multi-Commodity Mincost Flows

In this section, we consider multi-commodity mincost flows. In this setting, a solution consists of multiple flows satisfying different demands in the same networks. The flows do not mix but interact via the edge cost functions because the cost of an edge depends on the total flow of all commodity flows on that edge. Multi-commodity mincost flow are, for example, of interest in the computation of equilibria in traffic networks. In general, there is more than one population of drivers with different origins and destinations in the network. These populations interact via congestion (i.e., via costs) on the edges, but the flows cannot be mixed.

Formally, in the multi-commodity setting, we are given a set $K = \{1, \dots, k\}$ of $k \geq 1$ commodities and a demand vector $\mathbf{q}_j = (q_{v,j})_{v \in V} \in \mathbb{R}^n$ with $\sum_{v \in V} q_{v,j} = 0$ for all $k \in K$. We denote by $\mathbf{Q} = (q_{v,j})_{v \in V, j \in K} \in \mathbb{R}^{n \times k}$ the matrix of all demands, where the jth column is exactly the demand vectors of the commodity j.

A multi-commodity flow for demands $\mathbf{Q}$ is a matrix $\mathbf{X} = (x_{e,j})_{e \in E, j \in K} \in \mathbb{R}^{m \times k}$ satisfying the matrix equality $\boldsymbol{\Gamma}\mathbf{X} = \mathbf{Q}$, that is, flow conservation $\sum_{e \in \delta^-(v)} x_{e,j} - \sum_{e \in \delta^+(v)} x_{e,j} = q_{v,j}$ for all vertices $v \in V$ and all commodities $j \in K$. The columns of $\mathbf{X}$ are flow vectors $\mathbf{x}_j$ for demands $\mathbf{q}_j$ and, in particular, in the special case of $k = 1$, the matrix $\mathbf{X}$ is just a single-commodity flow as in the previous sections.

For every edge $e \in E$, we denote by $z_e := \sum_{j=1}^{k} x_{e,j}$ the total flow on that edge and by $\mathbf{z} = (z_e)_{e \in E} \in \mathbb{R}^m$ the vector of all total flows. In this section, we assume that the marginal cost functions f_e are piecewise linear and continuous. Then, $\mathbf{X} = (x_{e,j})_{e \in E, j \in K}$ is a multi-commodity mincost flow if it is a solution to

$$\min \sum_{e \in E} F_e\left(\sum_{j=1}^{k} x_{e,j}\right) \quad \text{s.t.} \quad \boldsymbol{\Gamma}\mathbf{X} = \mathbf{Q}, \quad \mathbf{X} \geq 0. \tag{20}$$

Similar to the single-commodity case, we obtain an optimality condition based on vertex potentials for the multi-commodity mincost problem.

Lemma 10. *A matrix $\mathbf{X} \in \mathbb{R}^{m \times k}$ is a multi-commodity minimum cost flow solving (20) if and only if there is a matrix $\boldsymbol{\Pi} \in \mathbb{R}^{n \times k}$ whose columns are vectors of vertex potentials $\boldsymbol{\pi}_j \in \mathbb{R}^n$ for every commodity $j \in K$ such that*

$$\begin{aligned}
\pi_{w,j} - \pi_{v,j} &= f_e(z_e) && \text{for all } e = (v,w) \in E \text{ with } x_{e,j} > 0, j \in K,\\
\pi_{w,j} - \pi_{v,j} &\leq f_e(z_e) && \text{for all } e = (v,w) \in E \text{ with } x_{e,j} = 0, j \in K,\\
\boldsymbol{\Gamma}\mathbf{X} &= \mathbf{Q}\\
\mathbf{X} &\geq 0,
\end{aligned}$$

where $z_e = \sum_{j=1}^{k} x_{e,j}$ is the total flow on edge $e \in E$.

The proof is analogous to the one of Lemma 1 and is thus omitted.

Because the objective function in (20) is not strictly convex, there is no unique multi-commodity mincost flow in general. This prevents a straightforward generalization of our method based on electrical networks for single-commodity networks, because there is no one-to-one correspondence between potentials and flows as in Lemma 2. Instead, we develop a similar method that computes *flow directions* $\Delta\mathbf{X}$ rather than potential directions $\Delta\boldsymbol{\pi}$.

Assume we want to solve the parametric variant of (20), that is, for all demands $\lambda\mathbf{Q}, \lambda \geq 0$, and that we are given a mincost flow $\mathbf{X}^0$ for some demand $\lambda^0\mathbf{Q}$. The basic idea of our multi-commodity variant is that the solution functions $\mathbf{x}(\lambda)$ are linear; that is, there is flow direction $\Delta\mathbf{X}$ such that $\mathbf{X}^0 + \lambda\Delta\mathbf{X}$ is a mincost flow for demands $(\lambda^0 + \lambda)$ for all $0 \leq \lambda \leq \lambda^*$. This direction $\Delta\mathbf{X}$ changes when either a breakpoint of an edge cost function is reached or an edge enters or leaves the support (i.e., an edge becomes used or unused). Our claim is that we can compute this direction by solving a convex program.

Given the mincost flow $\mathbf{X}^0$, let $\boldsymbol{\Pi}^0$ be the optimal dual vertex potentials from Lemma 10. Furthermore, let $\mathbf{z}^0$ denote the total flow with respect to $\mathbf{X}^0$. Then we define the sets

$$\begin{aligned}
E_j^{\mathrm{U}} &:= \{e \in E \mid x_{e,j}^0 > 0\}\\
E_j^{\mathrm{A}} &:= \{e \in E \mid \pi_{w,j}^0 - \pi_{v,j}^0 = F_e(z_e^0) \text{ and } x_{e,j}^0 = 0\}\\
E_j^{\mathrm{I}} &:= \{e \in E \mid \pi_{w,j}^0 - \pi_{v,j}^0 < F_e(z_e)\}
\end{aligned}$$

for every commodity $j \in K$. We call the edges in E_j^{U} *used edges* of commodity j. These edges carry a positive amount of flow, and thus, the flow on these edges may be increased or decreased when the overall demand changes. We refer to the edges in E_j^{A} as the *active but unused* edges. Flow on these edges can remain zero or increase when the demand changes. Finally, all other edges in E_j^{I} are considered *inactive edges* because they are not

of minimal cost and are thus not used when the demand increases or decreases by a small amount. Given a subset $F \subseteq E$ of the edges, let $\mathbb{1}_{e\in F} := 1$ if $e \in F$ and $\mathbb{1}_{e\in F} := 0$ if $e \notin F$ and let $\mathbf{I}_F := \operatorname{diag}(\mathbb{1}_{e_1\in F}, \dots, \mathbb{1}_{e_m\in F}) \in \mathbb{R}^m$. Furthermore, let $\mathbf{z}^+$ (respectively, $\mathbf{z}^-$) be the vector of positive (respectively, negative) parts of the total flows, that is, the vector with entries $z_e^+ := \max\{0, z_e\}$ (respectively, $z_e^- := \min\{0, z_e\}$). For every edge, define the values

$$a_e^+ := \lim_{z\downarrow z_e^0} \frac{f_e(z) - f_e(z_e^0)}{z - z_e^0} \quad \text{and} \quad a_e^- := \lim_{z\uparrow z_e^0} \frac{f_e(z) - f_e(z_e^0)}{z - z_e^0},$$

that is, the slopes of the marginal cost function f_e below and above z_e^0. By assumption, the marginal cost functions are piecewise linear; that is, they are of the form $f_e(z) = a_{e,t}z + b_{e,t}$. The values a_e^+ and a_e^- therefore coincide with the coefficients $a_{e,t}$. In particular, if z_e^0 is not a breakpoint of F_e, then a_e^+ and a_e^- coincide. Consider the convex quadratic program

$$\begin{aligned}
\min \quad & \frac{1}{2}\,\mathbf{z}^\top \mathbf{A}^+\,\mathbf{z}^+ + \frac{1}{2}\,\mathbf{z}^\top \mathbf{A}^-\,\mathbf{z}^- \\
\text{s.t.} \quad & \mathbf{z} = \sum_{j\in K} \mathbf{x}_j \\
& \mathbf{\Gamma}\,\mathbf{x}_j = \lambda^0 \mathbf{q}_j \quad \text{for all } j \in K, \\
& \mathbf{I}_{E_j^{\mathrm{A}}}\,\mathbf{x}_j \geq 0 \quad \text{for all } j \in K, \\
& \mathbf{I}_{E_j^{\mathrm{I}}}\,\mathbf{x}_j = 0 \quad \text{for all } j \in K.
\end{aligned} \tag{21}$$

Theorem 6. *Let the marginal cost functions f_e be piecewise linear, strictly increasing, and continuous and let $\mathbf{X}^0$ be a mincost flow for some demand $\lambda^0\mathbf{Q}$. Then,* (i) *the quadratic program* (21) *has an unique optimal solution $(\mathbf{X}^*, \mathbf{z}^*)$ that can be obtained in $\mathcal{O}(\mathrm{poly}(n, |K|))$-time.* (ii) *There is a $\lambda^* > 0$ such that for all $0 \leq \lambda \leq \lambda^*$ the flow $\mathbf{X}(\lambda) := \mathbf{X}^0 + \lambda\,\mathbf{X}^*$ is a multicommodity mincost flow for the demand $\mathbf{Q}(\lambda) = (\lambda^0 + \lambda)\,\mathbf{Q}$.*

Proof. By assumption, the objective function of (21) is strictly convex and bounded by zero from below. Furthermore, the feasible set is a closed convex set. Hence, an optimal solution exists and, because the objective function is strictly convex, it must be unique. The nonsmooth convex optimization problem can be solved with accuracy ϵ in $\mathcal{O}(\mathrm{poly}(n, |K|)\log(\frac{1}{\epsilon}))$ (Nesterov and Nemirovskii [47]) by the inscribed ellipsoid method (Khachiyan et al. [33]). Every optimal $\mathbf{X}^*$ is a solution to the KKT conditions that we discuss later. We obtain from these conditions that every optimal solution can be obtained as a solution of a system of linear equations. Cramer's rule allows to express solutions to this system as rational numbers whose denominators are the determinant of the involved matrix. Bounding this determinant with Hardamard's inequality, we obtain that every optimal solution is rational with a common denominator in $\mathcal{O}(M^R)$, where M is the largest entry of all matrices in (21) and $R \in \mathcal{O}(|K|n)$ is the number of linear constraints. Therefore, the exact optimal solution can be computed in $\mathcal{O}(\mathrm{poly}(n, |K|)\log(M))$ time.

Let $(\mathbf{X}^*, \mathbf{z}^*)$ be an optimal solution. Then the KKT conditions (Ruszczyński [51, theorem 3.34]) yield for $\mathbf{X} = \mathbf{X}^*$ and $\mathbf{z} = \mathbf{z}^*$ that

$$\begin{aligned}
\mathbf{0} \in\ & \partial_{\mathbf{X},\mathbf{z}} \frac{1}{2}\,\mathbf{z}^\top \mathbf{A}^+\,\mathbf{z}^+ + \partial_{\mathbf{X},\mathbf{z}} \frac{1}{2}\,\mathbf{z}^\top \mathbf{A}^-\,\mathbf{z}^- + \boldsymbol{\mu}^\top \partial_{\mathbf{X},\mathbf{z}} \left(\sum_{j\in K} \mathbf{x}_j - \mathbf{z} \right) \\
& + \sum_{j\in K} \boldsymbol{\phi}_j^\top \partial_{\mathbf{X},\mathbf{z}} (\lambda^0 \mathbf{Q} - \mathbf{\Gamma}\mathbf{x}_j) - \sum_{j\in K} \boldsymbol{\nu}_j^\top\, \partial_{\mathbf{X},\mathbf{z}}\, \mathbf{I}_{E_j^{\mathrm{A}}}\, \mathbf{x}_j + \sum_{j\in K} \boldsymbol{\eta}_j^\top\, \partial_{\mathbf{X},\mathbf{z}}\, \mathbf{I}_{E_j^{\mathrm{I}}}\, \mathbf{x}_j,
\end{aligned}$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\phi}_j, \boldsymbol{\nu}_j, \boldsymbol{\eta}_j$ for all $j \in K$ are vectors of dual variables with $\boldsymbol{\nu}_j \geq 0$. All subdifferentials with respect to $\mathbf{X}$ exclusively contain the gradient because the objective and all constraints are differentiable with respect to $\mathbf{X}$. The same holds for every subdifferential with respect to z_e as long as $z_e \neq 0$. Thus, for all edges e with $z_e^* \neq 0$, we get

$$\begin{aligned}
\mu_e &= a_e^+ z_e^* \quad \text{if } z_e^* > 0 \\
\text{and } \mu_e &= a_e^- z_e^* \quad \text{if } z_e^* < 0,
\end{aligned} \tag{22}$$

and, for all edges with $z_e^* = 0$, $\mu_e = \hat{a}_e z_e^* = 0$ for some $\hat{a}_e$ between a_e^- and a_e^+. The subdifferential with respect to $\mathbf{x}_j$ yields $\boldsymbol{\mu}^\top - \boldsymbol{\phi}_j^\top \mathbf{\Gamma} - \boldsymbol{\nu}_j^\top \mathbf{I}_{E_j^{\mathrm{A}}} + \boldsymbol{\eta}_j^\top \mathbf{I}_{E_j^{\mathrm{I}}} = \mathbf{0}$, which implies for every edge $e = (v, w) \in E$

$$\begin{aligned}
\mu_e &= \phi_{w,j} - \phi_{v,j} && \text{if } e \text{ is used,} \\
\mu_e &= \phi_{w,j} - \phi_{v,j} + \nu_{j,e} && \text{if } e \text{ is active but unused,} \\
\mu_e &= \phi_{w,j} - \phi_{v,j} - \eta_{j,e} && \text{if } e \text{ is inactive.}
\end{aligned}$$

Because $\nu_{j,e} \geq 0$ and $\nu_{j,e} = 0$ whenever $x^*_{e,j} \neq 0$ by the KKT conditions, we obtain

$$\begin{aligned} \mu_e &= \phi_{w,j} - \phi_{v,j} && \text{for all } e = (v,w) \text{ with } e \in E_j^{\mathrm{U}} \text{ or } e \in E_j^{\mathrm{A}} \text{ and } x^*_{e,j} > 0, \\ \mu_e &\geq \phi_{w,j} - \phi_{v,j} && \text{for all } e = (v,w) \text{ with } e \in E_j^{\mathrm{A}} \text{ and } x^*_{e,j} = 0. \end{aligned} \tag{23}$$

Thus, the flow $\mathbf{X}^*$ and the vertex potentials $\boldsymbol{\phi}_j$ satisfy the conditions from Lemma 10, ignoring inactive edges and admitting that the flow $\mathbf{X}^*$ can also be negative on edges $e \in E_j^{\mathrm{U}}$. In particular, every optimal solution is obtained as a solution of the active linear equations in (21), (22), and (23) and is therefore rational.

Define

$$\lambda_1^* := \min\left\{ -\frac{x^0_{e,j}}{x^*_{e,j}} \;\middle|\; e \in E, j \in K : x^*_{e,j} < 0 \right\},$$

$$\lambda_2^* := \min\left\{ \frac{F'_e(z^0_e) - (\pi^0_{w,j} - \pi^0_{v,j})}{\phi_{w,j} - \phi_{v,j}} \;\middle|\; e \in E_j^{\mathrm{I}}, j \in K : \phi_{w,j} - \phi_{v,j} > 0 \right\}.$$

The cost functions are piecewise functions with breakpoints $\tau_{e,t}$, $t = 0, \ldots, \bar{t}$. Denote by τ_e^- and τ_e^+ the closest breakpoints to z_e such that $z_e \in [\tau_e^-, \tau_e^+]$ and $z_e \neq \tau_e^-, \tau_e^+$ if $z_e > 0$. Note that τ_e^- may be zero and τ_e^+ may be ∞. Then define

$$\lambda_3^* := \min\left\{ -\frac{z^0_e - \tau^-_e}{z^*_e} \;\middle|\; e \in E : z^*_e < 0 \right\} \quad \text{and} \quad \lambda_4^* := \min\left\{ \frac{\tau^+_e - z^0_e}{z^*_e} \;\middle|\; e \in E : z^*_e > 0 \right\},$$

and let $\lambda^* := \min\{\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_4^*\}$. By definition, all λ_i^* values are positive (and may be infinite if no edge satisfies the conditions). It is easy to check that $\mathbf{X}^0 + \lambda \mathbf{X}^*$ is a flow satisfying demands $(\lambda^0 + \lambda)\mathbf{Q}$. The definition of λ_1^* ensures that for all $\lambda \leq \lambda^*$, all flows remain non-negative. The definitions of λ_3^* and λ_4^* ensure that no breakpoint of the cost functions is passed and the slope of $F'_e(z_e)$ is the same (namely, a_e^+ or a_e^-, respectively) for all $z_e = z^0_e + \lambda z^*_e$.

For every commodity $j \in K$, consider the vertex potential $\boldsymbol{\pi}_j(\lambda) := \boldsymbol{\pi}^0_j + \lambda\, \boldsymbol{\phi}_j$. Then, for every used edge $e \in E_j^{\mathrm{U}}$ and every active edge $e \in E_j^{\mathrm{A}}$ with $x^*_{e,j} > 0$, we obtain

$$\begin{aligned} \pi_{w,j}(\lambda) - \pi_{v,j}(\lambda) &= \pi^0_{w,j} - \pi^0_{v,j} + \lambda\,(\phi_{w,j} - \phi_{v,j}) \\ &\overset{(22)}{=} F'_e(z^0_e) + \lambda\,\mu_e \\ &\overset{(23)}{=} F'_e(z^0_e + \lambda\, z^*_e) = F'_e(z_e(\lambda)). \end{aligned}$$

Because $x_{e,j}(\lambda) > 0$ only if $e \in E_j^{\mathrm{U}}$ or $e \in E_j^{\mathrm{A}}$ and $x^*_{e,j} > 0$, the first part of the conditions in Lemma 10 is satisfied. If $e \in E_j^{\mathrm{A}}$ and $x^*_{e,j} = 0$, then we obtain similarly $\pi_{w,j}(\lambda) - \pi_{v,j}(\lambda) \leq F'_e(z_e(\lambda))$ because $\mu_e \geq \phi_{w,j} - \phi_{v,j}$ in this case. Finally, let $e \in E_j^{\mathrm{I}}$. Then $x^*_{e,j} = 0$ and thus $z_e(\lambda) = z^0_e$. Using the definition of λ_2^*, we obtain

$$\begin{aligned} \pi_{w,j}(\lambda) - \pi_{v,j}(\lambda) &= \pi^0_{w,j} - \pi^0_{v,j} + \lambda\,(\phi_{w,j} - \phi_{v,j}) \\ &\leq \pi^0_{w,j} - \pi^0_{v,j} + F'_e(z^0_e) - (\pi^0_{w,j} - \pi^0_{v,j}) = F_e(z_e(\lambda)), \end{aligned}$$

whenever $\phi_{w,j} - \phi_{v,j} > 0$. If $\phi_{w,j} - \phi_{v,j} < 0$, then $\pi_{w,j}(\lambda) - \pi_{v,j}(\lambda) < \pi^0_{w,j} - \pi^0_{v,j} < F'_e(z_e(\lambda))$. Thus, we confirmed the second part of the conditions in Lemma 10. Thus, all flows $\mathbf{X}(\lambda)$ satisfy the optimality conditions from Lemma 10 for $0 \leq \lambda \leq \lambda^*$ and hence the claim follows. □

Algorithm 2 (Multi-Commodity Algorithm)

Input: Directed graph G with piecewise linear, convex cost functions $F_e : \mathbb{R}_{\geq 0} \to \mathbb{R}$, demand matrix $\mathbf{Q}$
Output: Piecewise function $\mathbf{x} : \mathbb{R}_{\geq 0} \to \mathbb{R}^{k \times m}$ s.t. $\mathbf{x}(\lambda)$ is a mincost flow for all $\lambda \geq 0$
$\mathbf{X} \leftarrow 0$;
$\lambda^{\min} \leftarrow 0$;
repeat
  compute optimal solution $(\mathbf{X}^*, \mathbf{z}^*)$ of (21) wrt. $\mathbf{X}$;
  compute λ^*;
  $\lambda^{\max} \leftarrow \lambda^{\min} + \lambda^*$;
  $\mathbf{x}(\lambda) \leftarrow \mathbf{X} + \lambda \mathbf{X}^*$ for $\lambda \in [\lambda^{\min}, \lambda^{\max}]$;

$\mathbf{X} \leftarrow \mathbf{X} + \lambda^* \mathbf{X}^*$;
$\lambda^{\min} \leftarrow \lambda^{\max}$;
until $\lambda^{\max} = \infty$;
return x.

Theorem 6 gives rise to an algorithm to compute parametric multi-commodity mincost flows summarized in Algorithm 2. The algorithm works as follows: For $\lambda = 0$, start with the zero flow $\mathbf{X} = \mathbf{0}$. Then, in every iteration set up the quadratic program (21). In order to do so, we need to compute the optimal potentials $\mathbf{\Pi}$. These can be obtained by a shortest path computation with respect to the edges lengths $F_e'(0)$ for every commodity $j \in K$. Using these potentials, we can set up the edge sets $E_j^{\mathrm{U}}, E_j^{\mathrm{A}}$, and E_j^{I} for every commodity j and thus the quadratic program (21). With the solution $\mathbf{X}^*$ of the quadratic program, we can compute the value λ^* from Theorem 6—the latter also requires the dual variables $\phi_{v,j}$ that can be obtained as shortest path potentials w.r.t. the edge length $a_e z_e^*$. Then, Theorem 6 ensures that we can extend the solution function $\mathbf{x}$ up to $\lambda + \lambda^*$. Iterating this procedure then yields the piecewise linear solution function $\mathbf{x}$.

Lemma 11. *For cost functions F_e with finitely many breakpoints, Algorithm 2 terminates after finitely many iterations. Furthermore, every iteration takes $\mathcal{O}(\mathrm{poly}(n, |K|))$ time.*

Proof. We assume that there are only finitely many edges and finitely many breakpoints. Thus, there are only finitely many configurations of the sets $E_j^{\mathrm{U}}, E_j^{\mathrm{A}}$, and E_j^{I}, as well as there are only finitely many possible choices for the values a_e^+ and a_e^-. Suppose Algorithm 2 does not terminate in finitely many iterations. Then there is one configuration of these edge sets and a_e^+ and a_e^- occurring at least twice. Let us assume this configuration occurs for λ^1 and the flow $\mathbf{X}^1$ with potentials $\mathbf{\Pi}^1$, as well for $\lambda^2 > \lambda^1 + \lambda^{*,1}$ and the flow $\mathbf{X}^2$ with potentials $\mathbf{\Pi}^2$. Then the difference of these flows $\mathbf{X}^2 - \mathbf{X}^1$ together with the potential $\boldsymbol{\pi}^2 - \boldsymbol{\pi}^1$ satisfies the KKT conditions of the quadratic program (21). Because (21) has a unique solution $\mathbf{X}^*$, we have that $\mathbf{X}^2 - \mathbf{X}^1 = \mathbf{X}^*$. The definition of the values λ_1^* and λ_2^* ensure that all flows $\mathbf{X}^1 + \lambda \mathbf{X}^*$ are no mincost flows for $\lambda > \lambda^1 + \lambda^{1,*}$, implying that $\mathbf{X}^2$ is not a mincost flow, which is a contradiction. □

6. Complexity of Computing Parametric Mincost Flows

In this section, we discuss the complexity of our algorithms for computing electrical flows and mincost flows in the single- and multi-commodity case. For all variants of our algorithm (basic algorithm for electrical flows, extended algorithm for mincost flows, and the multi-commodity variant), each individual iteration of algorithm requires polynomial runtime. Furthermore, assuming nondegeneracy of the instances, every iteration produces one function part of the output flow functions $\lambda \mapsto \mathbf{x}(\lambda)$. The runtime of our algorithms depends on the complexity of the output flow functions. These can have a number of breakpoints that are exponentially in input size; that is, the size of the graph and number of breakpoints of the input functions in general (see Section 6.1.2). On a positive note, for series-parallel graphs, the number of breakpoints of the output functions is bounded by a polynomial of the input size, guaranteeing a polynomial runtime of our algorithm in this case (see Section 6.1.1).

6.1. Complexity for Single-Commodity Networks

We begin by discussing the complexity of our parametric algorithm for single commodities.

Theorem 7. *Let G be a graph on n vertices with homogeneous, strictly increasing, not necessarily continuous, marginal cost functions f_e, and let L be the number of breakpoints of the output flow functions. Then the solution curve and the output flow functions can be computed in* (i) $\mathcal{O}(n^{2.375} + n^2 L)$ *whenever the instance is nondegenerate and* (ii) $\mathcal{O}(n^{2.375} + \mathrm{poly}(n)L)$ *whenever the instance is degenerate.*

Proof. Consider Algorithm 1. The algorithm starts with the zero flow, sets up the initial Laplacian matrix in, and computes the first (generalized) inverse $\mathbf{L}_{\mathbf{t}}^*$. In the case of mincost flow (in contrast to electrical flows), this step requires the computation of an initial potential. Starting with $\boldsymbol{\pi} = \mathbf{0}$, Theorem 4 yields an initial potential in a nonambiguous region in $\mathcal{O}(n^2)$ time. The bottleneck of these initial steps is thus the computation of the inverse, which can be reduced to matrix multiplication (Corment et al. [13, chapter 28.4]). Using fast matrix multiplication, for example, the Coppersmith-Winograd-Algorithm (Coppersmith and Winograd [12]), the complexity of the initialization is $\mathcal{O}(n^{2.375})$.

In every iteration, the algorithm computes the values $\underline{\lambda}_{\mathbf{t}}$ and $\overline{\lambda}_{\mathbf{t}}$ for all edges. Because the vector $\boldsymbol{\gamma}_e$ contains only two nonzero elements, each of these values can be computed in constant time, and thus, the computation

and finding the minimum of all these values can be done in $\mathcal{O}(n^2)$ time. Finally, in every iteration, the matrix $\mathbf{L}_\mathbf{t}^*$ has to be updated. Using again the fact that $\boldsymbol{\gamma}_e$ contains only constantly many nonzero elements, the update formula from Theorem 2(i) implies a runtime of $\mathcal{O}(n^2)$ for every iteration.

In the nondegenerate case, every iteration produces exactly one breakpoint of the output functions. Thus, claim (i) follows. In the degenerate case, however, it is possible that algorithm needs multiple iterations before $\lambda^{\max}$ increases again. Thus, not all iterations correspond to breakpoints of the output function. In this case, we can use the quadratic program (21) to obtain a flow direction and a potential direction $\Delta\boldsymbol{\pi}$ (the dual variables ϕ in the proof of Theorem 6 correspond to the potential direction $\Delta\boldsymbol{\pi}$), proving (ii). □

If the cost are nonhomogeneous, that is, $f_e^-(0) \le 0 \le f_e^+(0)$ is not satisfied, the parametrized solution can still be obtained in output polynomial time. In this case, the minimal λ for which the problem is feasible must be determined by solving the linear program described in Section 4.3. Then, the initial solution and potential can be computed by solving the mincost flow problem for fixed demand $\lambda^{\min}\mathbf{q}$, which is possible in polynomial time, (e.g., with the algorithm described in Végh [56]).

In the nondegenerate case the complexity of the algorithm is determined by the complexity of solving the system $\mathbf{L}_\mathbf{t}\boldsymbol{\pi} = \mathbf{y}$ in every iteration. Initially, the matrix $\mathbf{L}_\mathbf{t}$ has to be inverted, resulting in the $\mathcal{O}(n^{2.375})$ term originating from the complexity of matrix multiplication. For all further iterations, the inverse can be updated with the formula from Theorem 2, resulting in the $\mathcal{O}(n^2)$ complexity for all further iterations. In general, a Laplacian system of the form $\mathbf{L}_\mathbf{t}\boldsymbol{\pi} = \mathbf{y}$ can be solved approximately much faster than in $\mathcal{O}(n^{2.375})$ time. The work of Spielman and Teng [54, 55] started a line of research (Cohen et al. [11], Kelner et al. [32], Kyng and Sachdeva [39]) establishing algorithms that find an ϵ-approximate solution to $\mathbf{L}_\mathbf{t}\boldsymbol{\pi} = \mathbf{y}$ in $\mathcal{O}(m\log^c n\log(\frac{1}{\epsilon}))$ time. Using these algorithms can improve the overall runtime of our algorithm significantly but only produce approximate solutions. To obtain an exact solution of $\mathbf{L}_\mathbf{t}\boldsymbol{\pi} = \mathbf{y}$ on rational input, ϵ has to be chosen in the order of $\mathcal{O}(\frac{1}{M^n})$, because the common denominator of the rational solution $\boldsymbol{\pi}$ is in $\mathcal{O}(M^n)$, where M is the largest absolute entry of the Laplacian matrix $\mathbf{L}_\mathbf{t}$, as argued in the proof of Theorem 6. Hence, using the fast Laplacian solvers have no better worst case complexity in this case.

Finally, we note that the explicit computation of the (generalized) inverse $\mathbf{L}_\mathbf{t}^*$ and its update via the formula of Theorem 2 may be numerically unstable. Rather than computing the inverse, we can also use the Cholesky decomposition of the positive definite submatrix $\hat{\mathbf{L}}_\mathbf{t}$. As for the inverse, there exists a rank 1 update formula for the Cholesky decomposition of a matrix (Gill et al. [20]) that can be computed in $\mathcal{O}(n^2)$ as well.

Both in the degenerate and nondegenerate cases, the complexity of the computation of the mincost flow functions depends on the number of breakpoints L of the mincost flow functions. We proceed to address the question whether L is also polynomial in the input size, that is, polynomial in the size of the graph and the number of breakpoints of the input functions. As we will show, in the special case of series-parallel networks, we can guarantee an input-polynomial runtime. In general, however, there are networks where the equilibrium flow functions have exponentially many breakpoints in the input size.

6.1.1. Series-Parallel Graphs. Following the definition of Duffin [17], we call two edges e, e' *confluent*, if there are not two (undirected) cycles $\mathcal{C}_1, \mathcal{C}_2$ in the graph G such that e, e' have the same orientation in $\mathcal{C}_1$ and different orientations in $\mathcal{C}_2$. Furthermore, we say an edge e is *s-t confluent* if the edge e and the additional virtual edge (t, s) are confluent.

Let $\mathbf{c} = (c_{e_1}, \ldots, c_{e_m})^\top \in \mathbb{R}^m_{\ge 0}$ be a vector of non-negative conductance values. We say a flow $\mathbf{x} \in \mathbb{R}^m$ is *induced by the potential* $\boldsymbol{\pi} \in \mathbb{R}^m$ with respect to $\mathbf{c}$, if $\mathbf{x} = \mathbf{C}\boldsymbol{\Gamma}^\top\boldsymbol{\pi}$, where $\mathbf{C} = \operatorname{diag}(c_{e_1}, \ldots, c_{e_m})$.

Theorem 8. *An edge e is s-t confluent if and only if the orientation of the flow on x_e of any s-t flow $\mathbf{x}$ that is induced by a potential $\boldsymbol{\pi}$ is independent of the conductance values $\mathbf{c}$, that is, either $x_e \ge 0$ or $x_e \le 0$ for all s-t flow $\mathbf{x}$ of any rate r and all conductance values $\mathbf{c}$.*

Proof. Because the edge e and the (virtual) edge (t, s) are confluent, the statement follows from Duffin [17, theorem 0]. □

Theorem 8 shows that the orientation of flow on *s-t* confluent edges is determined by the graph structure alone and does not depend on the conductance value. We proceed by showing that that series-parallel graphs are characterized by the confluence property.

We say a graph G is a two-terminal graph with terminals s and t if every edge e is contained in some *s-t* path. A two-terminal network with terminals s and t is *series-parallel* if it can be constructed by starting with the single edge (s, t) and repeatedly applying the following operations:

S: Replace an edge $e = (v, w)$ by two edges $e' = (v, u)$ and $e'' = (u, w)$ with a new vertex u.
P: Replace an edge $e = (v, w)$ by two edges $e' = (v, w)$ and $e'' = (v, w)$.

Theorem 9. *A two-terminal graph G with terminals s and t is series-parallel if and only if all edges are s-t confluent.*

Proof. We claim that all edges in G are s-t confluent if and only if all pairs of edges $e, e' \in E \cup \{(t, s)\}$ are confluent. The if direction is obvious. For the only-if direction, assume there are two edges $e = (v, w), e' = (v', w')$ that are not confluent. Because G is a two-terminal graph, there are s-t paths using both edges e and e'. Without loss of generality, we can assume that there exist $P_1 = (s, x_1, v, w, x_2, t)$ and $P_2 = (s, x_3, v', w', x_4, t)$, where x_1, x_2, x_3, x_4 are some sequences of vertices. By the assumption that e, e' are not confluent, there must exists two cycles with $\mathcal{C}_1 = (v, w, x_5, v', w', x_6, v)$ and $\mathcal{C}_2 = (v, w, x_7, w', v', x_8, v)$, where, again, x_5, x_6, x_7, x_8 are some sequences of vertices. However, then we can create the cycles $\mathcal{C}_3 := (s, x_1, v, w, x_2, t, s)$ and $\mathcal{C}_4 := (s, x_3, v', x_5, w, v, x_6, w', x_4, t, s)$ that include e and the (virtual) edge (t, s) in different orientations, yielding a contradiction to the assumption that e is s-t confluent. Hence, all edges in G are s-t confluent if and only if all pairs of edges $e, e' \in E \cup \{(t, s)\}$ are confluent. The latter is equivalent to G being series-parallel by Duffin [17, theorem 2]. □

Corollary 4. *Given a series-parallel graph G with piecewise-linear cost functions F_e with $\bar{t}_e$ breakpoints, let $\bar{t} := \sum_{e \in E} \bar{t}_e$. Then all output flows $f_e(\lambda)$ are monotone in the demand rate parameter λ. In particular, the solution curve $\boldsymbol{\pi}(\lambda)$ and the flow functions $\mathbf{x}(\lambda)$ have at most $\bar{t}$ many breakpoints, and thus the mincost flow flow functions can be computed in $\mathcal{O}(n^{2.375} + \bar{t} n^2 m)$ time and hence in polynomial time.*

Proof. Theorems 8 and 9 imply that the flows $\Delta \mathbf{x}_\mathbf{t} := \mathbf{C}_\mathbf{t} \boldsymbol{\Gamma}^\top \Delta \boldsymbol{\pi}_\mathbf{t}$ have the same orientation on every edge independent of the values $\mathbf{c}$. Thus, independent of the region, the flows $\Delta \mathbf{x}_\mathbf{t}$ have the sign and thus the functions $\mathbf{x}(\lambda)$ are monotone in the parameter λ. This implies, that the flow functions can cross every breakpoint at most once; hence, every breakpoint can contribute to at most one boundary crossing, implying the in particular part. □

6.1.2. Nested Braess Graphs. In this section, we present a family of graphs $G_j^\mathsf{B} = (V_j, E_j), j \in \mathbb{N}$ and cost functions $(F_e^j)_{e \in E_j, j \in \mathbb{N}}$ such that

a. The graph G_j^B has linear size in j, in particular $|V_j| = 2j + 2$ and $|E_j| = 4j + 1$.
b. The cost functions are $F_e^j(x) = \frac{1}{2} a_e^j x^2 + b_e^j x$. In particular, the cost function have no breakpoints (except for the artificial breakpoint at $x = 0$, as we restrict to non-negative flows).
c. The mincost flow functions $\mathbf{x}^j(\lambda)$ have $\Omega(2^j)$ many breakpoints.

Our construction uses coefficients $a_e^j = 0$, which we normally do not allow as we assume that all cost functions are strictly convex. The same construction works for slopes $a_e^j = \epsilon$ for some sufficiently small ϵ, but for ease of exposition, we omit the ϵ slopes.

For every flow $\mathbf{x}$, we denote by $S(\mathbf{x}) := \{e \in E : x_e > 0\}$ the *support* of $\mathbf{x}$. It is obvious that, if $\mathbf{x}^1, \mathbf{x}^2$ have different supports, then the corresponding potentials $\boldsymbol{\pi}^1, \boldsymbol{\pi}^2$ must be in different feasible regions $R_{\mathbf{t}^1}, R_{\mathbf{t}^2}$. Because every feasible regions corresponds to a function part of the output flow functions, the number of different supports is a lower bound on the number of the breakpoints of the mincost flow function $\mathbf{x}(\lambda)$.

For $j \in \mathbb{N}$, let $V_j := \{s = v_0, v_1, \dots, v_{2j}, t = v_{2j+1}\}$ and $E_j := E_j^1 \cup E_j^2 \cup E_j^3 \cup \{(v_j, v_{j+1})\}$, where

$$E_j^1 := \{(v_i, v_{i+1}) \mid i \in \{0, \dots, 2j\}, i \neq j\},$$
$$E_j^2 := \{(v_i, v_{2j-i}) \mid i \in \{0, \dots, j-1\}\},$$
$$E_j^3 := \{(v_{i+1}, v_{2j+1-i}) \mid i \in \{0, \dots, j-1\}\}.$$

We call the graph $G_j^\mathsf{B} := (V_j, E_j)$ the jth *nested Braess graph*. Every edge $e \in E_j$ is equipped with the cost function $F_e^j(x) = \frac{1}{2} a_e^j x^2 + b_e^j x$ with coefficients

$$a_e^j = \begin{cases} 1 & \text{if } e \in E_j^1 \setminus \{(v_j, v_{j+1})\}, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad b_e^j = \begin{cases} 10^{j-1-i} & \text{if } e = (v_i, v_{2j-i}), i = 0, \dots, j-1 \text{ or} \\ & \text{if } e = (v_{i+1}, v_{2j+1-i}), i = 0, \dots, j-1, \\ 0 & \text{otherwise.} \end{cases}$$

We compute parametric mincost flows for demands $\lambda \mathbf{q}, \lambda \geq 0$, where $q_s = -1$, $q_t = 1$, and $q_{v_i} = 0$ for $i = 1, \dots, 2j$; that is, we compute parametric s-t flows in the respective graphs.

For $j = 1$, this construction yields the classical Braess graph from the well-known paradox in the realm of traffic networks (Braess [9]) with flow dependent travel time functions x and 1 (corresponding to the marginal cost f_e in

Figure 10. Two nested Braess graphs G_j^{B} for $j = 1$ and $j = 3$ with associated cost functions on the edges.

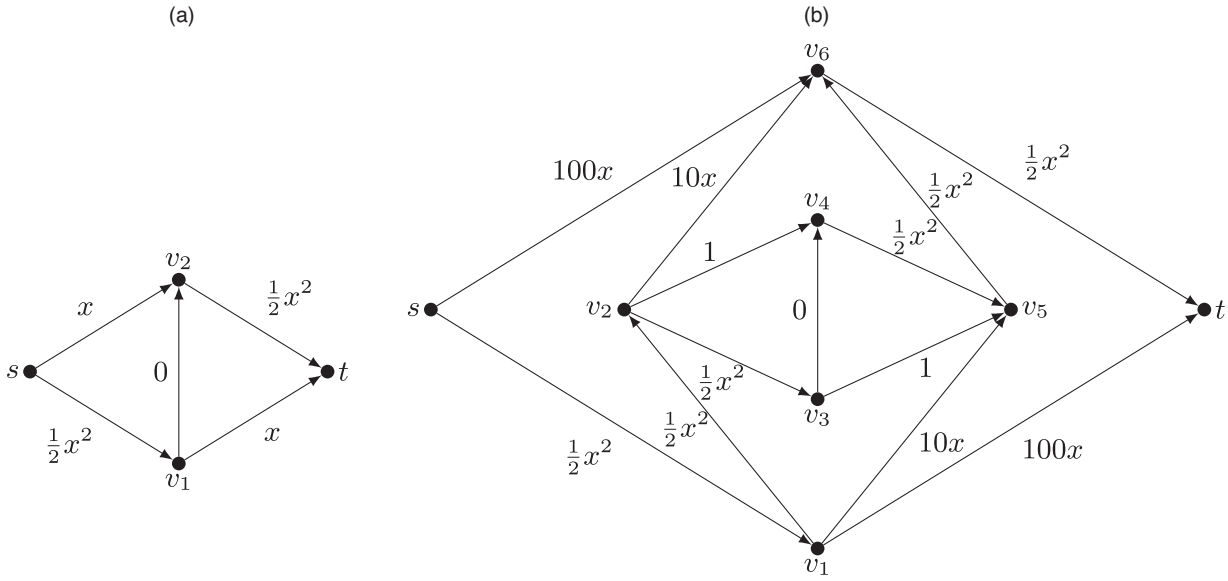

The first nested Braess graph G_1^{B}.

The third nested Braess graph G_3^{B}.

our notation) and 0 on the middle edge. Figure 10 depicts the construction for $j = 1$ and $j = 3$. We note that the subgraph $G_j^{\mathsf{B}}[V_j^p]$ induced by the subset of vertices $V_j^p := \{v_i : i \in \{p, \dots, 2j+1-p\}\}$ is again a nested Braess network, namely the $j-p$th nested Braess graph. Thus, $G_j^{\mathsf{B}}[V_j^p] = G_{j-p}^{\mathsf{B}}$ for $p < j$. By construction, the graphs G_j^{B} and the corresponding costs satisfy Properties (a) and (b). The following theorem proves the third claim (c).

Theorem 10. *For any $j \in \mathbb{N}$, consider the graph G_j^{B} with cost functions $(F_e^j)_{e\in E}$. Then there are 2^{j+1} different demand rates $\lambda_1, \dots, \lambda_{2^{j+1}} \in [0, 3 \cdot 10^{j-1})$ such that all corresponding support sets $S^{\lambda_1}, \dots, S^{\lambda_{2^{j+1}}}$ of the mincost flows are different. In particular, this means the mincost flow functions $\mathbf{x}(\lambda)$ have $\Omega(2^j)$ breakpoints.*

Proof. First, consider some fixed $j \in \mathbb{N}$. Let $\mathbf{x}(\lambda)$ be the mincost flow in G_j^{B} for the demand factor $\lambda \geq 0$. Denote by $\boldsymbol{\pi}(\lambda)$ the corresponding optimal potential from Lemma 1. (We assume without proof that both $\mathbf{x}(\lambda)$ and $\boldsymbol{\pi}(\lambda)$ are well-defined functions.) We claim that the mincost flow does not use the edges $e = (s, v_{2j})$ and $e = (v_1, t)$ for sufficiently small values λ, whereas for large values λ, the mincost flow uses the outermost edges $E_j^O := \{(s, v_1), (s, v_{2j}), (v_1, t), (v_{2j}, t)\}$ exclusively.

Claim 4. For fixed $j \in \mathbb{N}$, we have
(i) If $\lambda \geq 2 \cdot 10^{j-1}$, then $x_e(\lambda) = 0$ for all $e \notin E_j^O$ and $\pi_t(\lambda) - \pi_s(\lambda) = 10^{j-1} + \frac{\lambda}{2}$.
(ii) If $\lambda \leq 3 \cdot 10^{j-2}$, then $x_e(\lambda) = 0$ for $e = (s, v_{2j})$ and $e = (v_1, t)$.

Proof of Claim 4. For $\lambda \geq 2 \cdot 10^{j-1}$, consider the flow $\mathbf{x}(\lambda)$ with $x_e(\lambda) = \frac{\lambda}{2}$ for $e \in E_j^O$ and $x_e(\lambda) = 0$ otherwise. Define the vertex potential $\boldsymbol{\pi}(\lambda)$ by $\pi_s = 0$, $\pi_{v_{2j}} = 10^{j-1}$, $\pi_t = 10^{j-1} + \frac{\lambda}{2}$, and $\pi_v = \frac{\lambda}{2}$ otherwise. We claim, that $\boldsymbol{\pi}(\lambda)$ is an optimal potential for the flow $\mathbf{x}(\lambda)$ in the sense of Lemma 1. Indeed, for the edge $e = (s, v_1)$, we observe $f_e(x_e(\lambda)) = x_e(\lambda) = \frac{\lambda}{2} = \pi_{v_1} - \pi_s$. For $e = (v_1, t)$ we observe $f_e(x_e(\lambda)) = 10^{j-1} = \pi_t - \pi_{v_1}$. By symmetry, we get the corresponding equalities for the edges $e = (s, v_{2j})$ and $e = (v_{2j}, t)$. Thus, all outer flow-carrying edges satisfy the equality condition from Lemma 1. For the inner edges (with zero flow), we observe that the potential difference is zero except for the edges $e = (v_2, v_{2j})$ and $e = (v_{2j-1}, v_{2j})$ for which the potential difference is nonpositive. The latter follows from the assumption that implies $10^{j-1} - \frac{\lambda}{2} \leq 0$. Hence, all these edges satisfy the inequality condition from Lemma 1.

Now assume that $\lambda \leq 3 \cdot 10^{j-2}$ and $x_e(\lambda) > 0$ for $e = (s, v_{2j})$. (By symmetry, the following arguments hold likewise for $e = (v_1, t)$.) Using the second property from (i) for the inner component $G_j^{\mathsf{B}}[V_j^1] = G_{j-1}^{\mathsf{B}}$, we obtain that

$$\pi_{v_{2j}}(\lambda) - \pi_{v_1}(\lambda) \leq \pi_{v_{2j}}(3 \cdot 10^{j-2}) - \pi_{v_1}(3 \cdot 10^{j-2}) \leq 10^{j-2} + \frac{3 \cdot 10^{j-2}}{2}.$$

For the edge $e = (s, v_1)$, the optimality conditions imply $\pi_1(\lambda) - \pi_s(\lambda) \le f_e(x_e(\lambda)) = x_e(\lambda) \le 3 \cdot 10^{j-2}$. Therefore, we get for the edge $e = (s, v_{2j})$ that

$$\pi_{v_{2j}}(\lambda) - \pi_s(\lambda) = \pi_{v2j}(\lambda) - \pi_{v_1}(\lambda) + \pi_{v_1}(\lambda) - \pi_s(\lambda) \le 3 \cdot 10^{j-2} + \frac{3 \cdot 10^{j-2}}{2} < 10^{j-1} = f_e(x_e(\lambda)),$$

which is a contradiction to the optimality condition because we assumed $x_e(\lambda) > 0$, concluding the proof of the claim. $\triangle$

We prove the statement of the theorem by induction on j. For $j = 1$, we claim that $S^{\lambda_1} = \emptyset$ for $\lambda_1 = 0$, $S^{\lambda_2} = \{(v_0, v_1), (v_1, v_2), (v_2, v_3)\}$ for any $\lambda_2 \in (0, 1]$, $S^{\lambda_3} = E_1$ for any $\lambda_3 \in (1, 2)$, and $S^{\lambda_4} = E_1 \setminus \{(v_1, v_2)\}$ for any $\lambda_4 \in [2, 3]$. This may be verified easily by construction the corresponding optimal potentials. Hence, there are $4 = 2^{1+1}$ many different support sets for $j = 1$.

For $j \ge 2$, we obtain that for $\lambda \le 3 \cdot 10^{j-2}$, all flow is routed through the inner component $G_j^{\mathsf{B}}[V_j^1] = G_{j-1}^{\mathsf{B}}$ by Claim 4(ii). By induction hypotheses, there are 2^j different supports $\emptyset = \tilde{S}_1, \tilde{S}_2, \ldots, \tilde{S}_{2^j}$ for different demand rates $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_{2^j} \in [0, 3 \cdot 10^{j-2})$ in the inner component. Thus, the supports $S_1 = \emptyset$ and $S_i := \{(v_0, v_1), (v_{2j}, v_{2j+1})\} \cup \tilde{S}_i, i = 2, 3, \ldots, 2^j$ are supports for the mincost flows in G_j^{B} for demand rates $\lambda_1 = \tilde{\lambda}_1, \ldots, \lambda_{2^j} = \tilde{\lambda}_{2^j}$.

Claim 4(i) implies that for demand rates $\lambda \ge 2 \cdot 10^{j-1}$, all demand is routed over the outer paths in the network. Thus, the flow in the inner component decreases from at least $3 \cdot 10^{j-2}$ to 0. Because all flow functions are continuous, there must be demand rates $\lambda_{2^j+1}, \ldots, \lambda_{2^{j+1}}$ such that the flow in the inner component is exactly $\tilde{\lambda}_1, \ldots, \tilde{\lambda}_{2^j}$. At these demand rates, the overall mincost flow has the supports $S_{2^j+i} := \tilde{S}_i \cup \{(v_0, v_1), (v_0, v_{2j}), (v_1, v_{2j+1}), (v_{2j}, v_{2j+1})\}, i = 1, \ldots, 2^j$, that is, the supports of the inner component plus all outer edges. Overall, we have 2^{j+1} different supports. $\square$

6.2. Complexity for Multi-Commodity Networks

The multi-commodity case can be solved with Algorithm 2. Every iteration of the algorithm generates a new function part of the output mincost flow functions. Because every iteration takes polynomial time by Theorem 6, we obtain the following result.

Theorem 11. *Given a network with n vertices and $|K|$ commodities, denote by L the number of breakpoints of the output mincost flow functions. Then the solution curve and the mincost flow functions can be computed in $\mathcal{O}(\mathrm{poly}(n, |K|)L)$ time.*

Acknowledgment

These results were presented in preliminary form at the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (Klimm and Warode [35]).

Endnotes

[1] We ignore the first component for the lexicographic comparison because we want to compare the coefficient vectors $\mathbf{m}_{e,t}$ of the polynomials $p_{e,t}$. Because the vector $\boldsymbol{\varepsilon}$ has a zero as first component, the first coefficient is ignored and therefore not relevant to the comparison.

[2] The assumption $f_e^-(0) \le 0 \le f_e^+(0)$ implies that we only can model lower capacities $l_e \le 0$ and upper capacities $u_e \ge 0$. However, we show how to drop this assumption in the following section.

References

[1] Ahuja RK, Batra JL, Gupta SK (1984) A parametric algorithm for convex cost network flow and related problems. *Eur. J. Oper. Res.* 16: 222–235.
[2] Ahuja RK, Magnanti TL, Orlin JB (1993) *Network Flows: Theory, Algorithms, and Applications* (Prentice-Hall, Englewood Cliffs, NJ).
[3] Aissi H, Mahjoub AR, McCormick ST, Queyranne M (2015) Strongly polynomial bounds for multiobjective and parametric global minimum cuts in graphs and hypergraphs. *Math. Programming* 154(1–2):3–28.
[4] Arezki Y, van Vliet D (1990) A full analytical implementation of the Partan/Frank-Wolfe algorithm for equilibrium assignment. *Transportation Sci.* 24(1):58–62.
[5] Bar-Gera H (2002) Origin-based algorithm for the traffic assignment problem. *Transportation Sci.* 36(4):398–417.
[6] Beckmann MJ, McGuire CB, Winsten CB (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
[7] Best MJ (1996) An algorithm for the solution of the parametric quadratic programming problem. Fischer H, Riedmüller B, Schäffler S, eds. *Applied Mathematics and Parallel Computing* (Physica-Verlag).
[8] Birkhoff G, Diaz JB (1956) Non-linear network problems. *Quart. Appl. Math.* 13(4):431–443.
[9] Braess D (1968) Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12(1):258–268.

[10] Carstensen P (1983) The complexity of some problems in parametric linear and combinatorial programming. PhD thesis, University of Michigan, Ann Arbor.
[11] Cohen MB, Kyng R, Miller GL, Pachocki JW, Peng R, Rao AB, Xu SC (2014) Solving SDD linear systems in nearly $m\log^{1/2}n$ time. Shmoys D, ed. *Proc. 46th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 343–352.
[12] Coppersmith D, Winograd S (1987) Matrix multiplication via arithmetic progressions. Aho AV, ed. *Proc. 19th Annual ACM Sympos. Theory Comput* (Association for Computing Machinery, New York), 1–6.
[13] Cormen TH, Leiserson CE, Rivest RL, Stein C (2009) *Introduction to Algorithms*, 3rd ed. (MIT Press, Cambridge, MA).
[14] Correa J, Stier-Moses NE (2010) Wardrop equilibria. Cochran JJ, ed. *Wiley Encyclopedia of Operations Research and Management Science* (Wiley & Sons, Hoboken, NJ).
[15] Disser Y, Skutella M (2015) The simplex algorithm is NP-mighty. Indyk P, ed. *Proc. 26th Annual ACM-SIAM Sympos. Discrete Algorithms*, (Society for Industrial and Applied Mathematics, Philadelphia), 858–872.
[16] Duffin RJ (1947) Nonlinear networks. IIa. *Bull. Amer. Math. Soc. (New Series)* 53(10):963–971.
[17] Duffin RJ (1965) Topology of series-parallel networks. *J. Math. Anal. Appl.* 10(2):303–318.
[18] Frank M, Wolfe P (1956) An algorithm for quadratic programming. *Naval Res. Logist. Quart.* 3(1–2):95–110.
[19] Gallo G, Grigoriadis MD, Tarjan RE (1989) A fast parametric maximum flow algorithm and applications. *SIAM J. Comput.* 18(1):30–55.
[20] Gill PE, Golub GH, Murray W, Saunders MA (1974) Methods for modifying matrix factorizations. *Math. Comput.* 28(126):505–535.
[21] Goldberg PW, Papadimitriou CH, Savani R (2013) The complexity of the homotopy method, equilibrium selection, and Lemke-Howson solutions. *ACM Trans. Econom. Comput.* 1(2):1–25.
[22] Granot F, McCormick ST, Queyranne M, Tardella F (2012) Structural and algorithmic properties for parametric minimum cuts. *Math. Programming* 135(1):337–367.
[23] Grone R (1991) On the geometry and Laplacian of a graph. *Linear Algebra Appl.* 150:167–178.
[24] Gross M, Pfetsch ME, Schewe L, Schmidt M, Skutella M (2018) Algorithmic results for potential-based flows: Easy and hard cases. *Networks* 73(3):306–324.
[25] Hager WW (1989) Updating the inverse of a matrix. *SIAM Rev.* 31(2):221–239.
[26] Harks T, Kleinert I, Klimm M, Möhring RH (2015) Computing network tolls with support constraints. *Networks* 65(3):262–285.
[27] Harville DA (1997) *Matrix Algebra from a Statistician's Perspective* (Springer, New York).
[28] Herings PJJ, Peeters R (2010) Homotopy methods to compute equilibria in game theory. *Econom. Theory* 42(1):119–156.
[29] Herings PJJ, van den Elzen AH (2002) Computation of the Nash equilibrium selected by the tracing procedure in n-person games. *Games Econom. Behav.* 38(1):89–117.
[30] Karp RM, Orlin JB (1981) Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Appl. Math.* 3(1):37–45.
[31] Katzenelson J (1965) An algorithm for solving nonlinear resistor networks. *Bell Systems Tech. J.* 44(8):1605–1620.
[32] Kelner JA, Orecchia L, Sidford A, Allen Zhu Z (2013) A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. Boneh D, Roughgarden T, Feigenbaum J, eds. *Proc. 45th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 911–920.
[33] Khachiyan L, Tarasov S, Erlich E (1988) The inscribed ellipsoid method. *Soviet Math. Dokl.* 37(1):226–230.
[34] Kirchhoff G (1847) Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Vertheilung galvanischer Ströme geführt wird. *Ann. Phys.* 148(12):497–508.
[35] Klimm M, Warode P (2019) Computing all Wardrop equilibria parametrized by the flow demand. Chan TM, *Proc. 30th Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 917–934.
[36] Klimm M, Pfetsch M, Raber R, Skutella M (2020) On the robustness of potential-based flow networks. Technical report. Technische Universität Darmstadt, Darmstadt, Germany.
[37] Kojima M, Nishino H, Sekine T (1976) An extension of Lemke's method to the piece-wise linear complementarity problem. *SIAM J. Appl. Math.* 31(4):600–613.
[38] Kozlov MK, Tarasov SP, Khachiyan LG (1980) The polynomial solvability of convex quadratic programming. *USSR Comput. Math. Math. Phys.* 20(5):223–228.
[39] Kyng R, Sachdeva S (2016) Approximate Gaussian elimination for Laplacians—Fast, sparse, and simple. Ostrovsky R, Dinur I, eds. *Proc. 57th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society, Los Alamitos, CA), 573–582.
[40] LeBlanc LJ, Helgason RV, Boyce DE (1985) Improved efficiency of the Franke-Wolfe algorithm for convex network programs. *Transportation Sci.* 19(4):445–462.
[41] Lemke CE, Howson JT Jr (1964) Equilibrium points of bimatrix games. *SIAM J. Appl. Math.* 12(2):413–423.
[42] Merris R (1994) Laplacian matrices of graphs: A survey. *Linear Algebra Appl.* 197:143–176.
[43] Minoux M (1984) A polynomial algorithm for minimum quadratic cost flow problems. *Eur. J. Oper. Res.* 18(3):377–387.
[44] Minoux M (1986) Solving integer minimum cost flows with separable convex cost objective polynomially. Gallo G, Sandi C, eds. *Netflow at Pisa* (Springer, Berlin), 237–239.
[45] Mohar B, Alavi Y, Chartrand G, Oellermann OR (1991) The Laplacian spectrum of graphs. Alavi Y, Chartrand G, Ollermann O, Schwenk A, eds. *Graph Theory, Combinatorics, and Applications* (Wiley, New York), 871–898.
[46] Murty K (1980) Computational complexity of parametric linear programming. *Math. Programming* 19(1):213–219.
[47] Nesterov YE, Nemirovskii A (1994) *Interior-Point Polynomial Algorithms in Convex Programming* (Society of Industrial and Applied Mathematics, Philadelphia).
[48] Nikolova E, Kelner J, Brand M, Mitzenmacher M (2006) Stochastic shortest paths via quasi-convex maximization. Azar Y, Erlebach T, eds. *Proc. 14th Annual Eur. Sympos. Algorithms* (Springer, Berlin), 552–563.
[49] O'Hare SJ, Connors RD, Watling DP (2016) Mechanisms that govern how the Price of Anarchy varies with travel demand. *Transportation Res. Part B* 84:55–80.
[50] Rockafellar RT (1984) *Network Flows and Monotropic Optimization* (John Wiley and Sons, Hoboken, NJ).

[51] Ruszczyński A (2006) *Nonlinear Optimization* (Princeton University Press, Princeton, NJ).
[52] Sacher RS (1980) A decomposition algorithm for quadratic programming. *Math. Programming* 18(1):16–30.
[53] Sherman J, Morrison WJ (1950) Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Ann. Math. Statist.* 21(1):124–127.
[54] Spielman DA, Teng S (2004) Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. Láaszló Babai, ed. *Proc. 36th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 81–90.
[55] Spielman DA, Teng SH (2014) Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM J. Matrix Anal. Appl.* 35(3):835–885.
[56] Végh L (2016) A strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. *SIAM J. Comput.* 45(5):1729–1761.
[57] Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engrg.* 1(3):325–378.
[58] Youn H, Gastner MT, Jeong H (2008) Price of Anarchy in transportation networks: efficiency and optimality control. *Phys. Rev. Lett.* 101(12):128701-1–128701-4.
[59] Zadeh N (1973) A bad network problem for the simplex method and other minimum cost flow algorithms. *Math. Programming* 5(1): 255–266.