

Navigating Central Path with Electrical Flows: from Flows to Matchings, and Back

Aleksander Mądry*
EPFL
aleksander.madry@epfl.ch

Abstract

We present an $\tilde{O}(m^{\frac{10}{7}}) = \tilde{O}(m^{1.43})$ -time¹ algorithm for the maximum s - t flow and the minimum s - t cut problems in directed graphs with unit capacities. This is the first improvement over the sparse-graph case of the long-standing $O(m \min\{\sqrt{m}, n^{2/3}\})$ running time bound due to Even and Tarjan [ET75] and Karzanov [Kar73]. By well-known reductions, this also establishes an $\tilde{O}(m^{\frac{10}{7}})$ -time algorithm for the maximum-cardinality bipartite matching problem. That, in turn, gives an improvement over the celebrated $O(m\sqrt{n})$ running time bound of Hopcroft and Karp [HK73] and Karzanov [Kar73] whenever the input graph is sufficiently sparse.

At a very high level, our results stem from acquiring a deeper understanding of interior-point methods – a powerful tool in convex optimization – in the context of flow problems, as well as, utilizing certain interplay between maximum flows and bipartite matchings.

The core of our approach comprises a primal-dual algorithm for (near-)perfect bipartite \mathbf{b} -matching problem. This algorithm is inspired by path-following interior-point methods and employs electrical flow computations to gradually improve the quality of maintained solution by advancing it toward (near-)optimality along so-called central path. To analyze this process, we establish a formal connection that ties its convergence rate to the structure of corresponding electrical flows. Then, we exploit that connection to obtain a convergence guarantee for our algorithm that improves upon the well-known barrier of $\Omega(\sqrt{m})$ iterations corresponding to the generic worst-case performance bounds for interior-point-method-based algorithms. This improvement is based on refining certain insights into behavior of electrical flows that stem from the work of Christiano et al. [CKM⁺11] and combining them with a new technique for preconditioning primal-dual solutions.

The final ingredient of our approach is a simple reduction of the maximum s - t flow problem to the bipartite \mathbf{b} -matching problem. This reduction is then composed with the recent sub-linear-time algorithm for finding perfect matchings in regular graphs of Goel et al. [GKK10], to derive an efficient procedure for rounding fractional s - t flows and bipartite matchings.

*Part of this work was done when the author was with Microsoft Research New England.

¹We recall that $\tilde{O}(f)$ denotes $O(f \log^c f)$, for some constant c .

1 Introduction

The maximum s - t flow problem and its dual, the minimum s - t cut problem, are two of the most fundamental and extensively studied graph problems in combinatorial optimization [Sch03, AMO93, Sch02]. They have a wide range of applications (see [AMOR95]), are often used as subroutines in other algorithms (see, e.g., [AHK12, She09]), and a number of other important problems – e.g., bipartite matching problem [CLRS09] – can be reduced to them. Furthermore, these two problems were often a testbed for development of fundamental algorithmic tools and concepts. Most prominently, the Max-Flow Min-Cut theorem [EFS56, FF56] constitutes the prototypical primal-dual relation.

Several decades of extensive work resulted in a number of developments on these problems (see Goldberg and Rao [GR98] for an overview) and many of their generalizations and special cases. Still, despite all this effort, the basic problem of computing maximum s - t flow and minimum s - t cut in general graphs has resisted progress for a long time. In particular, the current best running time bound of $O(m \min\{m^{\frac{1}{2}}, n^{\frac{2}{3}}\} \log(n^2/m) \log U)$ (with U denoting the largest integer arc capacity) was established over 15 years ago in a breakthrough paper by Goldberg and Rao [GR98] and this bound, in turn, matches the $O(m \min\{m^{\frac{1}{2}}, n^{\frac{2}{3}}\})$ bound for unit-capacity graphs that Even and Tarjan [ET75] – and, independently, Karzanov [Kar73] – put forth over 35 years ago.

Recently, however, important progress was made in the context of undirected graphs. Christiano et al. [CKM⁺11] developed an algorithm that allows one to compute a $(1 + \varepsilon)$ -approximation to the undirected maximum s - t flow (and the minimum s - t cut) problem in $\tilde{O}(mn^{\frac{1}{3}}\varepsilon^{-11/3})$ time. Their result relies on devising a new approach to the problem that combines electrical flow computations with multiplicative weights update method (see [AHK12]). Later, Lee et al. [LRS13] presented a quite different – but still electrical-flow-based – algorithm that employs purely gradient-descent-type view to obtain an $\tilde{O}(mn^{1/3}\varepsilon^{-2/3})$ -time $(1 + \varepsilon)$ -approximation for the case of unit capacities. Finally, very recently, this line of work was culminated by Sherman [She13] and Kelner et al. [KLOS14] who independently showed how to integrate non-Euclidean gradient-descent methods with fast poly-logarithmic-approximation algorithms for cut problems of Mądry [Mađ10] to get an $O(m^{1+o(1)}\varepsilon^{-2})$ -time $(1 + \varepsilon)$ -approximation to the undirected maximum flow problem.

Finally, we note that, in parallel to the above work that is focused on designing weakly-polynomial algorithms for the maximum s - t flow and minimum s - t cut problems, there is also a considerable interest in obtaining running time bounds that are strongly-polynomial, i.e., that do not depend on the values of arc capacities. The current best such bound is $O(mn)$ and it follows by combining the algorithms of King et al. [KRT94] and Orlin [Orl13].

Bipartite Matching Problem. Another problem that we will be interested in is the (maximum-cardinality) bipartite matching problem – a fundamental assignment problem with numerous applications (see, e.g., [AMO93, LP86]) and long history, with its roots in the works of Frobenius [Fro12, Fro17] and König [Kön15, Kön16, Kön23] from the early 20th century (see [Sch05]). Already in 1931, König [Kön31] and Egerváry [Ege31] provided first constructive characterization of maximum matchings in bipartite graphs. This characterization can be turned into a polynomial-time algorithm. Then, in 1973, Hopcroft and Karp [HK73] – and, independently, Karzanov [Kar73] – devised the celebrated $O(m\sqrt{n})$ -time algorithm. Till date, this bound is the best one known in the regime of relatively sparse graphs. It can be improved, however, when the input graph is dense, i.e., when m is close to n^2 . In this case, one can combine the algebraic approach of Rabin

and Vazirani [RV89] – that itself builds on the work of Tutte [Tut47] and Lovász [Lov79] – with matrix-inversion techniques of Bunch and Hopcroft [BH74] to get an algorithm that runs in $O(n^\omega)$ time (see [Muc05]), where $\omega \leq 2.3727$ is the exponent of matrix multiplication [CW90, VW12]. Also, later on, Alt et al. [ABMP91], as well as, Feder and Motwani [FM95] developed combinatorial algorithms that offer a slight improvement – by a factor of, roughly, $\log_n \frac{n^2}{m}$ – over the $O(m\sqrt{n})$ bound of Hopcroft and Karp whenever the graph is sufficiently dense.

Finally, it is worth mentioning that there was also a lot of developments on the (maximum-cardinality) matching problem in general, i.e., not necessarily bipartite, graphs. Starting with the pioneering work of Edmonds [Edm65], these developments led to bounds that essentially match the running time guarantees that were previously known only for bipartite case. More specifically, the running time bound of $O(m\sqrt{n})$ for the general-graph case was obtained by Micali and Vazirani [MV80, Vaz94] (see also [GT91] and [GK04]). While, building on the algebraic characterization of the problem due to Rabin and Vazirani [RV89], Mucha and Sankowski [MS04] and then Harvey [Har09] gave $O(n^\omega)$ -time algorithms for general graphs.

1.1 Our Contribution

In this paper, we develop a new algorithm for solving maximum s - t flow and minimum s - t cut problems in directed graphs. More precisely, we prove the following theorem.

Theorem 1.1. *Let $G = (V, E)$ be a directed graph with m arcs and unit capacities. For any two vertices s and t , one can compute an integral maximum s - t flow and minimum s - t cut of G in $\tilde{O}(m^{\frac{10}{7}})$ time.*

This improves over the long-standing $O(m \min\{\sqrt{m}, n^{2/3}\})$ running time bound due to Even and Tarjan [ET75] and, in particular, finally breaks the $\Omega(n^{\frac{3}{2}})$ running time barrier for sparse directed graphs.

Furthermore, by applying a well-known reduction (see [CLRS09]), our new algorithm gives the first improvement on the sparse-graph case of the seminal $O(m\sqrt{n})$ -time algorithms of Hopcroft-Karp [HK73] and Karzanov [Kar73] for the maximum-cardinality bipartite matching problem.

Theorem 1.2. *Let $G = (V, E)$ be an undirected bipartite graph with m edges, one can solve the maximum-cardinality bipartite matching problem in G in $\tilde{O}(m^{\frac{10}{7}})$ time.*

This, again, breaks the 40-year-old running time barrier of $\Omega(n^{\frac{3}{2}})$ for this problem in sparse graphs.

Additionally, we design a simple reduction of the maximum s - t flow problem to perfect bipartite \mathbf{b} -matching problem (see Theorem 3.1). (This reduction can be seen as an adaptation of the reduction of the maximum vertex-disjoint s - t -path problem to the bipartite matching problem due to Hoffman [Hof60] – cf. Section 16.7c in [Sch03].²) As the reduction in the other direction is well-known already, this establishes an algorithmic equivalence of these two problems. We also show (see Theorem 3.3 and Corollary 3.4) how this reduction, together with the sub-linear-time algorithm for perfect matching problem in regular bipartite graphs of Goel et al. [GKK10], leads to an efficient, nearly-linear time, rounding procedure for s - t flows.³

²We thank Lap Chi Lau [Lau13] for pointing out this similarity.

³Recently, it came to our attention that a very similar rounding result was independently obtained by Khanna et al. [KKL13].

Finally, our main technical contribution is a primal-dual algorithm for (near-)perfect bipartite \mathbf{b} -matching problem (see Theorem 3.2). This iterative algorithm draws on ideas underlying interior-point methods and the electrical flow framework of Christiano et al. [CKM⁺11]. It employs electrical flow computations to gradually improve the quality of maintained solution by advancing it toward (near-)optimality along so-called central path.

We develop a way of analyzing this algorithm’s rate of convergence by relating it to the structure of the corresponding electrical flows (see Theorem 5.5). This understanding enables us to devise a way of perturbing (see Section 6.1) and preconditioning (see Section 6.2) our intermediate solutions to ensure a convergence in only $\tilde{O}(m^{\frac{3}{7}})$ iterations and thus improve over the well-known barrier of $\Omega(m^{\frac{1}{2}})$ iterations that all the previous interior-point-methods-based algorithms suffer from. (To the best of our knowledge, this is the first time that this barrier was broken for a natural optimization problem.)

We also note that most of this understanding of convergence behavior of interior-point methods can be carried over to general LP setting. Therefore, we are hopeful that our techniques can be extended and will eventually lead to breaking the $\Omega(m^{\frac{1}{2}})$ iterations barrier for general interior-point methods.

1.2 Our Approach

The core of our approach comprises two components. One of them is combinatorial in nature and exploits an intimate connection between the maximum s - t flow problem and bipartite matching problem. The other one is more linear-algebraic and relies on interplay of interior-point methods and electrical flows.

Maximum flows and bipartite matchings. The combinatorial component shows that not only one can reduce bipartite matching problem to the maximum s - t flow problem, but also that a reduction in the other direction exists. Namely, one can reduce, in a simple and purely combinatorial way, the maximum s - t flow problem to a certain variant of bipartite matching problem (see Theorem 3.1). Once this reduction is established, it allows us to shift our attention to the matching problem.

Also, as a byproduct, this reduction – together with the algorithm of Goel et al. [GKK10] – yields a fast procedure for rounding fractional maximum flows (see Corollary 3.4). This enables us to focus on obtaining solutions that are only nearly-optimal, instead of being optimal.

Bipartite Matchings and Electrical Flows. The other component is based on using the interior-point method framework in conjunction with nearly-linear time electrical flow computations, to develop a faster algorithm for the bipartite matching problem.

The point of start here is a realization that the recent approaches to approximating undirected maximum flow [CKM⁺11, LRS13, She13, KLOS13], despite achieving impressive progress, have fundamental limitations that make them unlikely to yield improvements for the exact undirected or (approximate) directed setting.⁴ Very roughly speaking, these limitations stem from the fact that, at their core, all these algorithms employ some version of gradient-descent method that relies on purely primal arguments, while almost completely neglecting the dual aspect of the problem. It is well-understood, however, that getting a running time guarantee that depends logarithmically,

⁴Note that it is known – see, e.g., [Mađ11] – that computing exact maximum s - t flow in undirected graphs is algorithmically equivalent to computing the exact or approximate maximum s - t flow in directed graph.

instead of polynomially, on ε^{-1} – and such dependence is a prerequisite to making progress in directed setting – one needs to also embrace the dual side of the problem and take full advantage of it.

Interior-point methods and fast algorithms. The above realization motivates us to consider a more sophisticated approach, one that is inherently primal-dual and achieves logarithmic dependence on ε^{-1} : interior-point methods. These methods constitute a powerful optimization paradigm that is a cornerstone of convex optimization (see, e.g., [BV04, Wri97, Ye97]) and already led to development of polynomial-time exact algorithms for a variety of problems. Unfortunately, despite all its advantages and successes in tackling hard optimization tasks, this paradigm has certain shortcomings in the context of designing fast algorithms. The main reason for that is the fact that each iteration of interior-point method requires solving of a linear system, a task for which the current fastest general-purpose algorithm runs in $O(n^\omega) = O(n^{2.3727})$ time [AHU74, CW90, VW12]. So, this bound becomes a bottleneck if one was aiming for, say, even sub-quadratic-time algorithm.

Fortunately, it turns out that there is a way to circumvent this issue. Namely, even though the above bound is the best one known in general, one can get a better running time when dealing with some specific problem. This is achieved by exploiting the special structure of the corresponding linear systems. A prominent (and most important from our point of view) example here is the family of flow problems. Daich and Spielman [DS08] showed that in the context of flow problems one can use the power of fast (approximate) Laplacian system solvers [?, KMP10, KMP11, KOSZ13] to solve the corresponding linear systems in nearly-linear time. This enabled [DS08] to develop a host of $\tilde{O}(m^{\frac{3}{2}})$ -time algorithms for a number of important generalizations of the maximum flow problem for which there was no such algorithms before.

Unfortunately, this bound of $\tilde{O}(m^{\frac{3}{2}})$ time turns out to also be a barrier if one wants to obtain even faster algorithms. The new difficulty here is that the best worst-case bound on the number of iterations needed for an interior-point method to converge to near-optimal solution is $\Omega(m^{1/2})$. Although it is widely believed that this bound is far from optimal, it seems that our theoretical understanding of interior-point method convergence is still insufficient to make any progress on this front. In fact, improving this state of affairs is a major and long-standing challenge in mathematical programming.

Beyond the $\Omega(m^{\frac{1}{2}})$ barrier. Our approach to circumventing this $\Omega(m^{\frac{1}{2}})$ barrier and obtaining the desired $\tilde{O}(m^{\frac{10}{7}})$ -time algorithm for the bipartite \mathbf{b} -matching problem consists of two stages.

First one – presented in Section 5 – corresponds to setting up a primal-dual framework for solving the near-perfect \mathbf{b} -matching problem. This framework is directly inspired by the principles underlying path-following interior-point methods and, in some sense, is equivalent to them. In it, we start with some initial sub-optimal solution (that is encoded as a minimum-cost flow problem instance) and gradually improve its quality up to near-optimality. These gradual improvements are guided by certain electrical flow computations – the flows are used to update the primal solution and the corresponding voltages update the dual one – and our solution ends up following a special trajectory in the feasible space: so-called central path.

We analyze the performance of this optimization process by establishing a formal connection that ties the size of each improvement step to a certain characteristic of the corresponding electrical flow. Very roughly speaking, this size (and thus the resulting rate of convergence) is directly related to how much the electrical flow we compute resembles the current primal solution (which is also a

flow). Once this connection is established, a simple energy-based argument immediately recovers the generic $O(m^{\frac{1}{2}})$ iterations bound known for interior-point methods. So, as each electrical flow computation can be performed in $\tilde{O}(m)$ time, this gives an overall $\tilde{O}(m^{\frac{3}{2}})$ -time algorithm.

Finally, to improve upon the above $O(m^{\frac{1}{2}})$ iterations bound and deliver the desired $O(m^{\frac{10}{7}})$ -time procedure, in Section 6, we devise two techniques: perturbation of arcs – that can be seen as a refinement of the edge removal technique of Christiano et al. [CKM⁺11]; and solution preconditioning – a way of adding auxiliary arcs to the solution to improve its conductance properties. We show that by a careful composition of these techniques, one is able to ensure that the guiding electrical flows align better with the primal solution – thus allowing taking larger progress steps and guaranteeing faster convergence – while keeping the unwanted impact of these modifications on the quality of final solution minimal. The analysis of this process constitutes the technical core of our result and is based on understanding of the interplay between the interior-point method and both the primal and dual structure of electrical flows.

We believe that this approach of understanding interior-point methods through the lens of electrical flows is a promising direction and our result is just a first step towards realizing its full potential.

1.3 Organization

We begin the technical part of the paper in Section 2 where we present some preliminaries on maximum flow problem, electrical flows, and bipartite (**b**-)matching problem, as well as, introduce some theorems we will need in the sequel. In Section 3, we provide a general outline of our results and the structure of our proof.

In Section 4, we describe the reduction of maximum *s-t* flow problem to the bipartite **b**-matching problem. Next, in Sections 5 and 6, we explain how our path-following algorithms and electrical flows can be used to get an improved algorithm for the bipartite **b**-matching problem, with Section 7 presenting the analysis of our path-following primitive. Finally, we conclude in Section 8 by showing how to round fractional **b**-matchings to integral ones.

2 Preliminaries

In this section, we introduce some basic notation and definitions we will need later.

2.1 σ -Flows and the Maximum *s-t* Flow Problem

Throughout this paper, we denote by $G = (V, E, \mathbf{u})$ a directed graph with vertex set V , arc set E (we allow parallel arcs), and (non-negative) integer capacities u_e , for each arc $e \in E$. We usually define $m = |E|$ to be the number of arcs of the graph in question and $n = |V|$ to be the number of its vertices. Each arc e of G is an ordered pair (u, v) , where u is its *tail* and v is its *head*.

The basic notion of this paper is the notion of a σ -flow in G , where $\sigma \in \mathbb{R}^n$, with $\sum_v \sigma_v = 0$, is the *demand vector*. By a σ -flow in G we understand any vector $\mathbf{f} \in \mathbb{R}^m$ that assigns values to arcs G and satisfies the *flow conservation constraints*:

$$\sum_{e \in E^+(v)} f_e - \sum_{e \in E^-(v)} f_e = \sigma_v, \quad \text{for each vertex } v \in V. \quad (1)$$

Here, $E^+(v)$ (resp. $E^-(v)$) is the set of arcs of G that are leaving (resp. entering) vertex v . Intuitively, these constraints enforce that the net balance of the total in-flow into vertex v and the total out-flow out of that vertex is equal to σ_v , for every $v \in V$.

Furthermore, we say that a σ -flow \mathbf{f} is *feasible* in G iff \mathbf{f} obeys the *non-negativity and capacity constraints*:

$$0 \leq f_e \leq u_e, \quad \text{for each arc } e \in E. \quad (2)$$

One type of σ -flows that will be of special interest to us are s - t flows, where s (the *source*) and t (the *sink*) are two distinguish vertices of G . Formally, a σ -flow \mathbf{f} is an s - t flow iff its demand vector σ is equal to $F \cdot \chi_{s,t}$ for some $F \geq 0$ – we call F the *value* of \mathbf{f} – and the demand vector $\chi_{s,t}$ that has -1 (resp. 1) at the coordinate corresponding to s (resp. t) and zeros everywhere else.

Now, the *maximum s - t flow problem* corresponds to a task of finding for a given graph $G = (V, E, \mathbf{u})$, a source s , and a sink t , a feasible s - t flow \mathbf{f}^* in G of maximum value F . We call such a flow \mathbf{f}^* that maximizes F the *maximum s - t flow of G* and denote its value by F^* .

Sometimes, we will be also interested in (uncapacitated) *minimum-cost σ -flow problem* (with non-negative costs). In this problem, we have a directed graph G with infinite capacities on arcs (i.e., $u_e = +\infty$, for all e) and certain (*non-negative*) *length* (or *cost*) l_e assigned to each arc e . Our goal is to find a feasible σ -flow \mathbf{f} in G whose *cost* $l(\mathbf{f}) := \sum_e l_e f_e$ is minimal. (Note that as we have infinite capacities here, the feasibility constraint (2) just requires that $f_e \geq 0$ for all arcs e .)

Finally, one more problem that will be relevant in this context is the *minimum s - t cut problem*. In this problem, we are given a directed graph $G = (V, E, \mathbf{u})$ with integer capacities, as well as, a source s and sink t , and our task is to find an s - t cut $C \subseteq V$ in G minimizes the *capacity* $\mathbf{u}(C) := \sum_{E^-(C)} u_e$ among all s - t cuts. Here, a cut $C \subseteq V$ is an s - t cut iff $s \in C$ and $t \notin C$, and $E^-(C)$ is the set of all arcs (u, v) with $u \in C$ and $v \notin C$. It is well-known [EFS56, FF56] that the minimum s - t cut problem is the dual of the maximum s - t problem and, in particular, that the capacity of the minimum s - t cut is equal to the value of the maximum s - t flow, as well as, that given a maximum s - t flow one can easily obtain the corresponding minimum s - t cut.

2.2 Undirected Graphs

Although the focus of our results is on directed graphs, it will be crucial for us to consider undirected graphs too. To this end, we view an undirected graph $G = (V, E, \mathbf{u})$ as a directed one in which the ordered pair $(u, v) \in E$ does not denote an arc anymore, but an (undirected) *edge* (u, v) and the order just specifies an *orientation* of that edge from u to v . (Even though we use the same notation for these two different types of graphs, we will always make sure that it is clear from the context whether we deal with directed graph that has arcs, or with undirected graph that has edges.) From this perspective, the definitions of σ -flow \mathbf{f} that we introduced above for directed graphs transfer over to undirected setting almost immediately. The only (but very crucial) difference is that in undirected graphs a feasible flow can have some of f_e s being negative - this corresponds to the flow flowing in the direction that is opposite to the edge orientation. As a result, the feasibility condition (2) becomes

$$|f_e| \leq u_e, \quad \text{for each arc } e \in E. \quad (3)$$

Also, the set $E^+(v)$ (resp. $E^-(v)$) denotes now the set of incident edges that are oriented towards (resp. away) from v , and $E(v) := E^+(v) \cup E^-(v)$ is just the set of all edges incident to v , regardless of their orientation.

Finally, given a directed graph $G = (V, E, \mathbf{u})$, by its *projection* \bar{G} we understand an undirected graph that arises from treating each arc of G as an edge with the corresponding orientation. Note that if G had two arcs (u, v) and (v, u) then \bar{G} will have two parallel edges (u, v) and (v, u) that have opposite orientation and, possibly, different capacities.

2.3 Electrical Flows and Potentials

A notion that will play a fundamental role in this paper is the notion of electrical flows. Here, we just briefly review some of the key properties that we will need later. For an in-depth treatment we refer the reader to [Bol98].

Consider an undirected graph G and some vector of resistances $\mathbf{r} \in \mathbb{R}^m$ that assigns to each edge e its *resistance* $r_e > 0$. For a given σ -flow \mathbf{f} in G , let us define its *energy* (with respect to \mathbf{r}) $\mathcal{E}_r(\mathbf{f})$ to be

$$\mathcal{E}_r(\mathbf{f}) := \sum_e r_e f_e^2 = \mathbf{f}^T \mathbf{R} \mathbf{f}, \quad (4)$$

where \mathbf{R} is an $m \times m$ diagonal matrix with $R_{e,e} = r_e$, for each edge e .

For a given undirected graph G , a demand vector σ , and a vector of resistances \mathbf{r} , we define an *electrical σ -flow* in G (that is *determined* by resistances \mathbf{r}) to be the σ -flow that minimizes the energy $\mathcal{E}_r(\mathbf{f})$ among all σ -flows in G . As energy is a strictly convex function, one can easily see that such a flow is unique. Also, we emphasize that we do *not* require here that this flow is feasible with respect to capacities of G (cf. (3)). Furthermore, whenever we consider electrical flows in the context of a directed graph G , we will mean an electrical flow – as defined above – in the (undirected) projection \bar{G} of G .

One of very useful properties of electrical flows is that it can be characterized in terms of vertex potentials inducing it. Namely, one can show that a σ -flow \mathbf{f} in G is an electrical σ -flow determined by resistances \mathbf{r} iff there exist *vertex potentials* ϕ_v (that we collect into a vector $\phi \in \mathbb{R}^n$) such that, for any edge $e = (u, v)$ in G that is oriented from u to v ,

$$f_e = \frac{\phi_v - \phi_u}{r_e}. \quad (5)$$

In other words, a σ -flow \mathbf{f} is an electrical σ -flow iff it is *induced* via (5) by some vertex potential ϕ . (Note that orientation of edges matters in this definition.)

Using vertex potentials, we are able to express the energy $\mathcal{E}_r(\mathbf{f})$ (see (4)) of an electrical σ -flow \mathbf{f} in terms of the potentials ϕ inducing it as

$$\mathcal{E}_r(\mathbf{f}) = \sum_{e=(u,v)} \frac{(\phi_v - \phi_u)^2}{r_e}. \quad (6)$$

One of the consequences of this characterization of electrical flows via vertex potentials is that one can view the energy of an electrical σ -flow as being a result of optimization not over all the σ -flows but rather over certain set of vertex potentials. Namely, we have the following lemma that, for completeness, we prove in the Appendix A.

Lemma 2.1. *For any graph $G = (V, E)$, any vector of resistances \mathbf{r} , and any demand vector σ ,*

$$\frac{1}{\mathcal{E}_r(\mathbf{f}^*)} = \min_{\phi | \sigma^T \phi = 1} \sum_{e=(u,v) \in E} \frac{(\phi_v - \phi_u)^2}{r_e},$$

where \mathbf{f}^* is the electrical σ -flow determined by \mathbf{r} in G . Furthermore, if ϕ^* are the vertex potentials corresponding to \mathbf{f}^* then the minimum is attained by taking ϕ to be equal to $\tilde{\phi} := \phi^*/\mathcal{E}_{\mathbf{r}}(\mathbf{f}^*)$.

Note that the above lemma provides a convenient way of lowerbounding the energy of an electrical σ -flow. One just needs to expose any vertex potentials ϕ such that $\sigma^T \phi = 1$ and this will immediately constitute an energy lowerbound. Also, another basic but useful property of electrical σ -flows is captured by the following fact.

Fact 2.2 (Rayleigh Monotonicity). *For any graph $G = (V, E)$, demand vector σ and any two vectors of resistances \mathbf{r} and \mathbf{r}' such that $r_e \geq r'_e$, for all $e \in E$, we have that if \mathbf{f} (resp. \mathbf{f}') is the electrical σ -flow determined by \mathbf{r} (resp. \mathbf{r}') then*

$$\mathcal{E}_{\mathbf{r}}(\mathbf{f}) \geq \mathcal{E}_{\mathbf{r}'}(\mathbf{f}').$$

2.4 Laplacian Solvers

A very important algorithmic property of electrical flows is that one can compute very good approximations of them in nearly-linear time. Below, we briefly describe the tools enabling that.

To this end, let us recall that electrical σ -flow is the (unique) σ -flow induced by vertex potentials via (5). So, finding such a flow boils down to computing the corresponding vertex potentials ϕ . It turns out that computing these potentials can be cast as a task of solving certain type of linear system called *Laplacian* systems. To see that, let us define the *edge-vertex incidence matrix* \mathbf{B} being an $n \times m$ matrix with rows indexed by vertices and columns indexed by edges such that

$$\mathbf{B}_{v,e} = \begin{cases} 1 & \text{if } e \in E^+(v), \\ -1 & \text{if } e \in E^-(v), \\ 0 & \text{otherwise.} \end{cases}$$

Now, we can compactly express the flow conservation constraints (1) of a σ -flow \mathbf{f} (that we view as a vector in \mathbb{R}^m) as

$$\mathbf{B}\mathbf{f} = \sigma.$$

On the other hand, if ϕ are some vertex potentials, the corresponding flow \mathbf{f} induced by ϕ via (5) (with respect to resistances \mathbf{r}) can be written as

$$\mathbf{f} = \mathbf{R}^{-1} \mathbf{B}^T \phi,$$

where again \mathbf{R} is a diagonal $m \times m$ matrix with $R_{e,e} := r_e$, for each edge e .

Putting the two above equations together, we get that the vertex potentials ϕ that induce the electrical σ -flow determined by resistances \mathbf{r} are given by a solution to the following linear system

$$\mathbf{B}\mathbf{R}^{-1} \mathbf{B}^T \phi = \mathbf{L}\phi = \sigma, \tag{7}$$

where $\mathbf{L} := \mathbf{B}\mathbf{R}^{-1} \mathbf{B}^T$ is the (weighted) *Laplacian* \mathbf{L} of G (with respect to the resistances \mathbf{r}). One can easily check that \mathbf{L} is an $n \times n$ matrix indexed by vertices of G with entries given by

$$L_{u,v} = \begin{cases} \sum_{e \in E(v)} 1/r_e & \text{if } u = v, \\ -1/r_e & \text{if } e = (u, v) \in E, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

One can see that the Laplacian \mathbf{L} is not invertible, but – as long as, the underlying graph is connected – it’s null-space is one-dimensional and spanned by all-ones vector. As we require our demand vectors $\boldsymbol{\sigma}$ to have its entries sum up to zero (otherwise, no $\boldsymbol{\sigma}$ -flow can exist), this means that they are always orthogonal to that null-space. Therefore, the linear system (7) has always a solution $\boldsymbol{\phi}$ and one of these solutions⁵ is given by

$$\boldsymbol{\phi} = \mathbf{L}^\dagger \boldsymbol{\sigma},$$

where \mathbf{L}^\dagger is the Moore-Penrose pseudo-inverse of \mathbf{L} .

Now, from the algorithmic point of view, the crucial property of the Laplacian \mathbf{L} is that it is symmetric and *diagonally dominant*, i.e., for any $v \in V$, $\sum_{u \neq v} |L_{u,v}| \leq L_{v,v}$. This enables us to use fast approximate solvers for symmetric and diagonally dominant linear systems to compute an approximate electrical $\boldsymbol{\sigma}$ -flow. Namely, building on the work of Spielman and Teng [ST03, ST04], Koutis et al. [KMP10, KMP11] designed an SDD linear system solver that implies the following theorem. (See also recent work of Kelner et al. [KOSZ13] that presents an even simpler nearly-linear-time Laplacian solver.)

Theorem 2.3. *For any $\varepsilon > 0$, any graph G with n vertices and m edges, any demand vector $\boldsymbol{\sigma}$, and any resistances \mathbf{r} , one can compute in $\tilde{O}(m \log m \log \varepsilon^{-1})$ time vertex potentials $\tilde{\boldsymbol{\phi}}$ such that $\|\tilde{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_{\mathbf{L}} \leq \varepsilon \|\boldsymbol{\phi}^*\|_{\mathbf{L}}$, where \mathbf{L} is the Laplacian of G , $\boldsymbol{\phi}^*$ are potentials inducing the electrical $\boldsymbol{\sigma}$ -flow determined by resistances \mathbf{r} , and $\|\boldsymbol{\phi}\|_{\mathbf{L}} := \sqrt{\boldsymbol{\phi}^T \mathbf{L} \boldsymbol{\phi}}$.*

To understand the type of approximation offered by the above theorem, observe that $\|\boldsymbol{\phi}\|_{\mathbf{L}}^2 = \boldsymbol{\phi}^T \mathbf{L} \boldsymbol{\phi}$ is just the energy of the flow induced by vertex potentials $\boldsymbol{\phi}$. Therefore, $\|\tilde{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_{\mathbf{L}}$ is the energy of the electrical flow $\tilde{\mathbf{f}}$ that “corrects” the vertex demands of the electrical $\tilde{\boldsymbol{\sigma}}$ -flow induced by potentials $\tilde{\boldsymbol{\phi}}$, to the ones that are dictated by $\boldsymbol{\sigma}$. So, in other words, the above theorem tells us that we can quickly find an electrical $\tilde{\boldsymbol{\sigma}}$ -flow $\tilde{\mathbf{f}}$ in G such that $\tilde{\boldsymbol{\sigma}}$ is a slightly perturbed version of $\boldsymbol{\sigma}$ and $\tilde{\mathbf{f}}$ can be corrected to the electrical $\boldsymbol{\sigma}$ -flow \mathbf{f}^* that we are seeking, by adding to it some electrical flow $\tilde{\mathbf{f}}$ whose energy is at most ε fraction of the energy of the flow \mathbf{f}^* . (Note that electrical flows are linear, so we indeed have that $\mathbf{f}^* = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}$.) As we will see, this kind of approximation is completely sufficient for our purposes.

2.5 Bipartite \mathbf{b} -Matchings

A fundamental graph problem that constitutes both an application of our results, as well as, one of the tools we use to establish them, is the (*maximum-cardinality*) *bipartite \mathbf{b} -matching problem*. In this problem, we are given an undirected bipartite graph $G = (V, E)$ with $V = P \cup Q$ – where P and Q are the two sets of bipartition – as well as, a *demand vector* \mathbf{b} that assigns to every vertex v an integral and positive demand b_v . Our goal is to find a maximum cardinality *multiset* M of the edges of G that forms a *\mathbf{b} -matching*. That is, we want to find a multi-set M of edges of G that is of maximum cardinality subject to a constraint that, for each vertex $v \in V$, the number of edges of M that are incident to v is at most b_v . (When $b_v = 1$ for every vertex v , we will simply call such M a *matching*.)

We say that a \mathbf{b} -matching M is *perfect* iff every vertex in V has exactly b_v edges incident to it in M . Note that a perfect \mathbf{b} -matching – if it exists in G – has to necessarily be of maximum

⁵Note that the linear system (7) will have many solutions, but each two of them are equivalent up to a translation. So, as the formula (5) is translation-invariant, each of these solutions will yield the same unique electrical $\boldsymbol{\sigma}$ -flow.

cardinality. Also, if a graph has a perfect \mathbf{b} -matching then it must be that $\sum_{v \in P} b_v = \sum_{v \in Q} b_v$. Now, by the *perfect bipartite \mathbf{b} -matching problem* we mean a task in which we need to either find the perfect \mathbf{b} -matching in G or conclude that it does not exist.

Finally, by a *fractional* solution to a \mathbf{b} -matching problem, we understand an $|E|$ -dimensional vector \mathbf{x} that allocates non-negative value of x_e to each edge e and is such that for every vertex v of G , the sum $\sum_{e \in E(v)} x_e$ of (fractional) incident edges in \mathbf{x} is at most b_v . Also, we define the *size* of a fractional \mathbf{b} -matching \mathbf{x} to be $|\mathbf{x}|_1$.

An interesting class of graphs that is guaranteed to always have a perfect matching are bipartite graphs that are *d-regular*, i.e., that have the degree of each vertex equal to d . A remarkable algorithm of Goel et al. [GKK10] shows that one can find a perfect matching in such graphs in time that is proportional only to number of its vertices and not edges. (Note that a d -regular bipartite graph has exactly $\frac{dn}{2}$ edges and thus this number can be much higher than n when d is large.) In particular, they prove the following theorem that we will use later.

Theorem 2.4 (see Theorem 4 in [GKK10]). *Given an $n \times n$ doubly-stochastic matrix \mathbf{M} with m non-zero entries, one can find a perfect matching in the support of \mathbf{M} in $O(n \log^2 n)$ expected time with $O(m)$ preprocessing time.*

3 From Flows to Matchings, and Back

As we already mentioned, our results stem from exploiting the interplay between the maximum s - t flow and bipartite \mathbf{b} -matching problem, as well as, from understanding the performance of interior-point methods – when applied to these two problems – via the structure of corresponding electrical flows. To highlight these elements, we decompose the proof of our main theorem (Theorem 1.1) into three natural parts.

Reducing Maximum Flow to \mathbf{b} -Matching

First, we focus on analyzing the relationship between the maximum s - t flow and the (maximum-cardinality) bipartite \mathbf{b} -matching problem. It is well-known that the latter can be reduced to the former in a simple way. As it turns out, however, one can also go the other way – there is a simple, combinatorial reduction from the maximum flow problem to the task of finding a perfect bipartite \mathbf{b} -matching.⁶

Before making this precise, let us introduce one definition. Consider a \mathbf{b} -matching problem instance corresponding to a bipartite graph $G = (V, E)$ with P and Q ($V = P \cup Q$) being two sides of the bipartition. For any edge $e = (p, q) \in E$, let us define the *thickness* $d(e)$ of that edge to be $d(e) := \min\{b_p, b_q\}$. (So, $d(e)$ is an upper bound on the value of x_e in any feasible \mathbf{b} -matching \mathbf{x} .) We say that a \mathbf{b} -matching instance is *balanced* iff

$$\sum_{e \in E} d(e) \leq 4|\mathbf{b}|_1. \tag{9}$$

Now, in Section 4, we establish the following result.

⁶One can view this as one possible explanation of why the techniques used in the context of bipartite matchings and maximum flows are so similar.

Theorem 3.1. *If one can solve a balanced instance of a perfect bipartite \mathbf{b} -matching problem in a (bipartite) graph with \bar{n} vertices and \bar{m} edges in $T(\bar{n}, \bar{m}, |\mathbf{b}|_1)$ time, then one can solve the maximum s - t flow problem in a graph $G = (V, E, \mathbf{u})$ with m arcs and capacity vector \mathbf{u} in $\tilde{O}((m + T(\Theta(m), 4m, 4|\mathbf{u}|_1)) \log |\mathbf{u}|_1)$ time.*

This connection between maximum flows and bipartite matchings is useful in two ways. Firstly, it enables us to reduce the main problem we want to solve – the maximum s - t flow problem with unit capacities – to a seemingly simpler one: the perfect bipartite \mathbf{b} -matching problem. Secondly, the fact that this reduction works also for fractional instances provides us with an ability to lift our \mathbf{b} -matching rounding procedure that we develop later (see Theorem 3.3) to the maximum flow setting (see Corollary 3.4).

The Algorithm for Near-Perfect \mathbf{b} -Matching Problem

Once the above reduction is established, we can proceed to designing an improved algorithm for the perfect bipartite \mathbf{b} -matching problem. This algorithm consists of two parts.

The first one – constituting the technical core of our paper – is related to the (fractional) near-perfect bipartite \mathbf{b} -matching problem, a certain relaxation of the perfect bipartite \mathbf{b} -matching problem. To describe this task formally, let us call a \mathbf{b} -matching \mathbf{x} *near-perfect* if its size $|\mathbf{x}|_1$ is at least $\frac{|\mathbf{b}|_1}{2} - \tilde{O}(m^{\frac{3}{7}})$, i.e., it is within $\tilde{O}(m^{\frac{3}{7}})$ additive factor of the size of a perfect \mathbf{b} -matching. Now, given a bipartite graph $G = (P \cup Q, E)$ and demand vector \mathbf{b} , the *near-perfect \mathbf{b} -matching problem* is a task of either finding a *near-perfect \mathbf{b} -matching* in G or concluding that no *perfect \mathbf{b} -matching* exists in that graph.

Our goal is to design an algorithm that solves this near-perfect \mathbf{b} -matching problem in $\tilde{O}(m^{\frac{10}{7}})$ time. To this end, in Sections 5 and 6 we prove the following theorem.

Theorem 3.2. *Let $G = (V, E)$ with $V = P \cup Q$ be an undirected bipartite graph with n vertices and m edges and let \mathbf{b} be a demand vector that corresponds to a balanced \mathbf{b} -matching instance with $|\mathbf{b}|_1 = O(m)$. In $\tilde{O}(m^{\frac{10}{7}})$ time, one can either find a fractional near-perfect \mathbf{b} -matching \mathbf{x} or conclude that no perfect \mathbf{b} -matching exists in G .*

(Observe that whenever we have an instance of maximum s - t flow problem that has \bar{m} arcs and unit capacities, $|\mathbf{u}|_1$ is exactly \bar{m} . So, if we apply the reduction from Theorem 3.1 to that instance then the resulting \mathbf{b} -matching problem instance will be balanced, have $m \leq 4\bar{m}$ edges, as well as, $|\mathbf{b}|_1 \leq 4|\mathbf{u}|_1 = 4\bar{m} \leq 2m$. Therefore, we will be able to apply the above Theorem 3.2 to it.)

At a very high level, our algorithm for the near-perfect \mathbf{b} -matching problem is inspired by the way the existing interior-point method path-following algorithms (see, e.g., [Ye97, Wri97, BV04]) can be used to solve it. Basically, our algorithm is an iterative method that starts with some initial, far-from-optimal solution and then gradually improves this maintained solution to near-optimality (pushing it along so-called central path) using appropriate electrical flows as a guidance. We then show how to tie the convergence rate of this process to the structure of the guiding electrical flows. At that point, one can use a simple energy-bounding argument to establish a generic convergence bound that yields an (unsatisfactory) $\tilde{O}(m^{\frac{3}{2}})$ -time algorithm.

To improve upon this bound and deliver the desired $\tilde{O}(m^{\frac{10}{7}})$ -time algorithm, we show how one can appropriately “shape” these guiding electrical flows to make their guidance more effective and thus guarantee faster convergence. Very roughly speaking, it turns out there is a way of changing

the maintained solution to make it essentially the same from the point of view of our \mathbf{b} -matching instance, while dramatically improving the quality of corresponding electrical flows that guide it.

Our way of executing this idea is based on a careful composition of two techniques. One of them corresponds to perturbing, in a certain way, the arcs that are most significantly distorting the structure of electrical flow – this technique can be viewed as a refinement of edge removal technique of Christiano et al. [CKM⁺11]. The other technique corresponds to preconditioning the whole solution by adding additional, auxiliary, arcs to it. These arcs are chosen so to significantly improve the conductance properties of the solution (when viewed as a graph with resistances) while not leading to too significant deformation of the final obtained solution.

Rounding Near-Perfect \mathbf{b} -Matchings

Finally, our final step on our way towards solving the perfect \mathbf{b} -matching problem (and thus the maximum s - t flow problem) is related to turning the approximate and fractional answer returned by the algorithm from Theorem 3.2 into an exact and integral one. To this end, note that if that algorithm returned a near-perfect \mathbf{b} -matching that was integral, there would be a standard way to either turn it into a perfect \mathbf{b} -matching or conclude that no such perfect \mathbf{b} -matching exists. Namely, one could just use repeated augmenting path computations. It is well-known that given an integral \mathbf{b} -matching, one can perform, in $O(m)$ time, an augmenting path computation that either results in increasing the size of our \mathbf{b} -matching by one, or concludes that no further augmentation is possible (and thus no perfect \mathbf{b} -matching exists). So, as our initial near-perfect \mathbf{b} -matching has size at least $\frac{|b|_1}{2} - \tilde{O}(m^{\frac{3}{7}})$, after at most $\tilde{O}(m^{\frac{3}{7}})$ iterations, i.e., in time $\tilde{O}(m^{\frac{10}{7}})$, we would get the desired answer.

Unfortunately, the above approach can fail completely once our near-perfect \mathbf{b} -matching is fractional. This is so, as in this case we do not have any meaningful lowerbound on the progress on the size of the \mathbf{b} -matching brought by the augmenting path computation.

Therefore, to deal with this issue, we develop the last ingredient of our algorithm: a nearly-linear time procedure that allows one to round fractional \mathbf{b} -matchings. More precisely, in Section 8, building on the work of Goel et al. [GKK10] (see Theorem 2.4), we establish the following theorem.

Theorem 3.3. *Let $G = (V, E)$ be an undirected bipartite graph with m edges and let \mathbf{b} be a demand vector, if \mathbf{x} is a fractional \mathbf{b} -matching in G of size $k = |\mathbf{x}|_1$ then one can find in $\tilde{O}(m)$ time an integral \mathbf{b} -matching in G of size $\lfloor k \rfloor$.*

Clearly, if we apply the above rounding method to the fractional near-perfect matching \mathbf{x} computed by the algorithm from Theorem 3.2, it will give us an integral \mathbf{b} -matching \mathbf{x}^* whose size is still at least $\frac{|b|_1}{2} - \tilde{O}(m^{\frac{3}{7}})$. So, the augmenting path-based approach we outlined above will let us obtain the desired integral and exact answer to the perfect \mathbf{b} -matching problem within the desired time bound.

In the light of all the above, we see that combining all the above pieces indeed yields an $\tilde{O}(m^{\frac{10}{7}})$ -time algorithm for the perfect bipartite \mathbf{b} -matching problem in graphs with $|b|_1 = O(m)$. Now, using the reduction from Theorem 3.1, this gives us the analogous algorithm for the maximum s - t flow problem in unit-capacity graphs and that, in turn, results in an algorithm for the bipartite matching problem. So, both Theorem 1.1 and Theorem 1.2 hold.

Rounding s - t Flows

Finally, we mention the other byproduct of our techniques – the fast rounding procedure for flows. Namely, using the reduction described in Theorem 3.1 and the rounding from Theorem 3.3 we can obtain a fast rounding procedure not only for fractional \mathbf{b} -matchings but also for fractional s - t flows. Specifically, the proof of the following corollary appears in Appendix B.

Corollary 3.4. *Let $G = (V, E, \mathbf{u})$ be a directed graph with capacities and let \mathbf{f} be some feasible fractional s - t flow in G of value F . In $\tilde{O}(m)$ time, we can obtain out of \mathbf{f} an integral s - t flow \mathbf{f}^* of value $\lfloor F \rfloor$ that is feasible in G .*

Again, we note that a very similar rounding result was independently obtained by Khanna et al. [KKL13].

4 From Maximum Flows to Perfect Matchings

In this section, we show how to reduce the maximum s - t flow problem in a directed capacitated graph $G = (V, E, \mathbf{u})$ to solving $O(\log |\mathbf{u}|_1)$ balanced instances of the perfect bipartite \mathbf{b} -matching problem, i.e., we prove Theorem 3.1. We note that our reduction can be seen as an adaptation of the reduction of the maximum vertex-disjoint s - t path problem to the bipartite matching problem due to Hoffman [Hof60] – cf. Section 16.7c in [Sch03].

To this end, let $G = (V, E, \mathbf{u})$ with $n = |V|$ vertices and $m = |E|$ arcs, as well as, the source s and sink t be our input instance of the maximum s - t flow problem. Without loss of generality, we can assume that there is no arcs entering s and no arcs leaving t , as these arcs do not affect the maximum s - t flow. Also, let F^* be the value of the maximum s - t flow in G .

4.1 The Reduction

We show that for any integral value of F , we can setup, in $\tilde{O}(m)$ time, a balanced bipartite \mathbf{b} -matching problem instance, for some demands \mathbf{b} and bipartite graph $\bar{G} = (P \cup Q, \bar{E})$, such that: (1) there will be a perfect \mathbf{b} -matching in \bar{G} if there is a feasible s - t flow of value F in G ; and (2) given a perfect \mathbf{b} -matching in \bar{G} one can recover in $\tilde{O}(m)$ time an s - t flow of value F that is feasible in G . Observe that once such a reduction is designed, Theorem 3.1 will follow by noticing that $1 \leq F^* \leq |\mathbf{u}|_1$ and applying a simple binary search strategy to find the value of F^* and extract the corresponding maximum s - t -flow.

Given the input graph $G = (V, E, \mathbf{u})$, source s , sink t and the value of F , the construction of our desired balanced bipartite \mathbf{b} -matching instance $\bar{G} = (P \cup Q, \bar{E})$ is as follows. First, for each arc $e \in E$, we create two vertices $p_e \in P$ and $q_e \in Q$ and an edge (p_e, q_e) between them, as well as, we set the demand b_{p_e} and b_{q_e} of these vertices to u_e . Next, for every vertex v of G other than s and t , we add a vertex p_v to P and a vertex q_v to Q . Also, we create an edge (p_v, q_v) , as well as, an edge (p_v, q_e) (resp. (q_v, p_e)) for every arc e that is incoming to (resp. outgoing of) v in G . We set the demands b_{p_v} (resp. b_{q_v}) to be equal to $\sum_{e \in E^+(v)} u_e$ (resp. $\sum_{e \in E^-(v)} u_e$). Finally, we create a vertex $q_s \in Q$ (resp. $p_t \in P$) and add an edge (q_s, p_e) (resp. (q_e, p_t) for each arc e that is leaving s (resp. incoming to t) in G . We put the demand b_{q_s} (resp. b_{p_t}) to be $(\sum_{e \in E^-(s)} u_e) - F$ (resp. $(\sum_{e \in E^+(t)} u_e) - F$). (Note that we can assume here that both these quantities are non-negative as both $\sum_{e \in E^-(s)} u_e$ and $\sum_{e \in E^+(t)} u_e$ are obvious upperbounds on the value of F^* .)

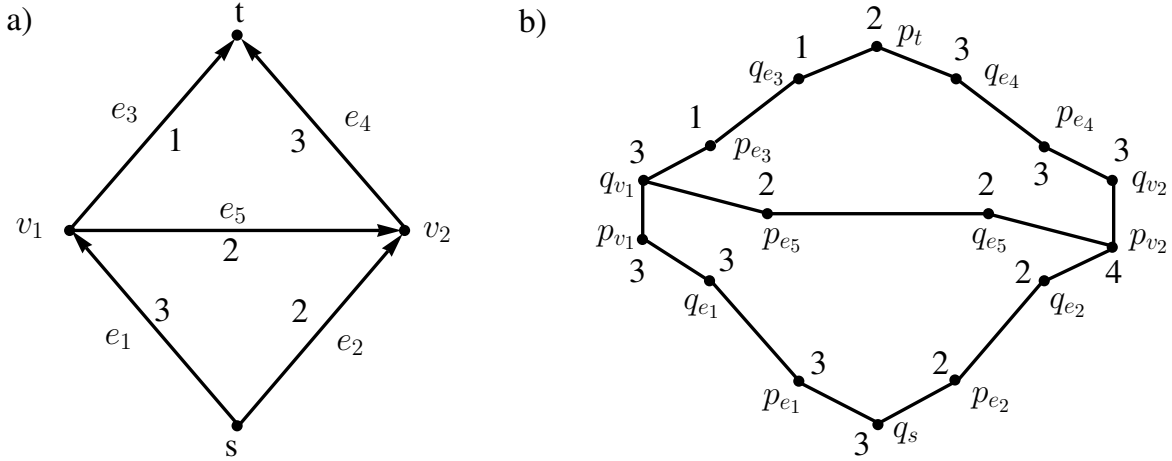


Figure 1: **a)** An example directed s - t flow instance G . Numbers next to arcs denote their capacities. **b)** The \mathbf{b} -matching instance corresponding to the example from a) in case of $F = 2$. Here, numbers next to vertices denote their demands.

An example s - t flow instance and the corresponding instance of the bipartite \mathbf{b} -matching can be found in Figure 1.

To see that this instance is balanced, note that every edge h of \bar{G} that is incident to some vertex p_e or q_e has its thickness $d(h)$ equal to $u_e = b_{p_e} = b_{q_e}$. So, the contribution of these edges to the total thickness $\sum_{h \in \bar{E}} d(h)$ of edges of \bar{G} is at most $3 \sum_{e \in E} u_e \leq \frac{3}{2} |\mathbf{b}|_1$. On the other hand, the only edges that are not incident to some p_e or q_e are the ones of the form (p_v, q_v) . However, the total contribution of these edges to the total thickness is at most

$$\sum_{v \neq s, t} \min \left\{ \sum_{e \in E^+(v)} u_e, \sum_{e \in E^-(v)} u_e \right\} \leq \sum_{v \neq s, t} \frac{\sum_{e \in E^+(v)} u_e + \sum_{e \in E^-(v)} u_e}{2} \leq |\mathbf{u}|_1 \leq |\mathbf{b}|_1,$$

as needed.

Now, the proof of correctness of this reduction appears in Appendix C.

5 Basic $\tilde{O}(m^{\frac{3}{2}})$ -Time Algorithm for Bipartite \mathbf{b} -Matching Problem

Over the next two sections, we prove Theorem 3.2. That is, we present an algorithm for the near-perfect bipartite \mathbf{b} -matching problem in the setting where the input instance is balanced (see (9)) and $|\mathbf{b}|_1$ is $O(m)$. In what follows we assume, for convenience, that $|\mathbf{b}|_1$ is at most $2m$ and that the graph G is sparse, i.e., $m = O(n)$.⁷

⁷It is easy to see that these assumptions are made without loss of generality. Whenever $|\mathbf{b}|_1$ is $O(m)$, one can ensure that $|\mathbf{b}|_1 \leq 2m$ and $m = O(n)$ by adding an appropriate – but still $O(m)$ – number of dummy copies of complete bipartite $K_{6,6}$ graph with uniform demands. Adding each such dummy isolated copy brings the ratio of $|\mathbf{b}|_1$ and m , as well as, of m to n down towards $\frac{18}{12}$, while never leading to violation of the balance condition (9) and preserving the \mathbf{b} -matching structure of the original input graph.

In this section, we show a basic algorithm that runs in $\tilde{O}(m^{\frac{3}{2}})$ time. Later, in Section 6, we refine this algorithm to obtain the desired running time of $\tilde{O}(m^{\frac{10}{7}})$.

For the sake of clarity, in our description and analysis we assume that the nearly-linear time Laplacian system solver (see Theorem 2.3) always returns an exact solution, i.e., all the electrical σ -flows we compute are exact. We discuss how to handle the approximate nature of the solver's output in Appendix E.9.

From b -Matching to Min-Cost σ -flow

Let us fix our instance of the bipartite b -matching problem in bipartite graph $G = (V, E)$ with $V = P \cup Q$. We will solve our b -matching instance by reducing it to a task of finding a minimum-cost $\hat{\sigma}$ -flow in a certain related directed graph $\hat{G} = (\hat{V}, \hat{E}, \hat{l})$ with \hat{l} being a length vector.

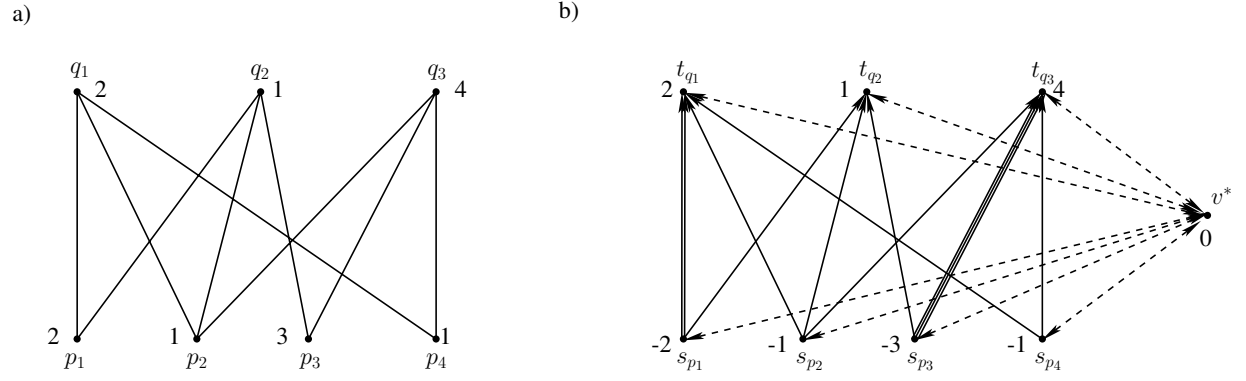


Figure 2: **a)** An example instance of bipartite b -matching problem. Numbers next to vertices represent their demands. **b)** The minimum-cost $\hat{\sigma}$ -flow problem instance corresponding to the example from a). All arcs have cost \hat{l}_e equal to 1 and the numbers next to vertices denote their demands in $\hat{\sigma}$. There are two parallel copies of the arc (s_{p_1}, t_{q_1}) and three parallel copies of the arc (s_{p_3}, t_{q_3}) . Also, each dashed arc represents two arcs that have the same endpoints but opposite orientation.

The reduction is performed as follows (see Figure 2 for an example). The vertex set \hat{V} of the graph \hat{G} consist of a special vertex v^* , as well as, vertices s_p (resp. t_q), for every vertex $p \in P$ (resp. $q \in Q$) of the graph G . Next, for every edge $e = (p, q)$ in G , we add to \hat{G} $d(e)$ copies of an arc (s_p, t_q) , where we recall that $d(e) := \min\{b_p, b_q\}$ is the thickness of e . Finally, for each vertex $p \in P$ (resp. $q \in Q$) of G , we add to \hat{G} arcs (s_p, v^*) and (v^*, s_p) (resp. (v^*, t_q) and (t_q, v^*)). We set the lengths \hat{l}_e of all arcs e to 1.

To gain some intuition on this reduction, note that if a perfect b -matching indeed exists in G then the flow that encodes it in \hat{G} is fully supported on the arcs (s_p, t_q) and does not send more than one unit of flow on any of these arcs. So, the purpose of including the extra vertex v^* and the arcs incident to it is to support (and appropriately penalize) the initial and intermediate solutions as they approach optimality.

Also, observe that this new graph has $\hat{n} := n + 1$ vertices and, due to our b -matching instance

being balanced, we have that the total number \widehat{m} of arcs is at most

$$2n + \sum_{e=(p,q) \in G} d(e) \leq 2n + O(m) = O(m).$$

So, bounding our running time in terms of \widehat{m} provides a bound in terms of the number of edges m of our original \mathbf{b} -matching instance that is asymptotically the same.

Now, consider a demand vector $\widehat{\sigma}$ that has surplus of b_p at each vertex s_p , a deficit of b_q at each vertex t_q and a zero demand at vertex v^* . (Note that such a demand vector will be valid, i.e., $\sum_v \widehat{\sigma}_v = 0$, as we can assume that $\sum_p b_p = \sum_q b_q$ – otherwise it would be impossible to have a perfect \mathbf{b} -matching in G .) We claim that any near-optimal $\widehat{\sigma}$ -flow gives us a solution to our near-perfect \mathbf{b} -matching instance. (Recall from Section 3 that a \mathbf{b} -matching is near-perfect if its size is at least $\frac{|b|_1}{2} - \widetilde{O}(\widehat{m}^{\frac{3}{7}})$. Although, in the lemma below it suffices that we have a slack of only $\frac{1}{2}$ instead of $\widetilde{O}(\widehat{m}^{\frac{3}{7}})$.)

Lemma 5.1. *Given any feasible $\widehat{\sigma}$ -flow \mathbf{f} in \widehat{G} whose cost $\widehat{\ell}(\mathbf{f})$ is within additive $\frac{1}{2}$ of the optimum, in $\widetilde{O}(\widehat{m})$ time, we can either compute a (fractional) near-perfect \mathbf{b} -matching \mathbf{x} in G or conclude that no perfect \mathbf{b} -matching exists in G .*

Proof. First, observe that if there exists a perfect \mathbf{b} -matching \mathbf{x}^* in G then a flow \mathbf{f}^* that just puts, for each $e = (p, q)$ of G , $\frac{x_e^*}{d(e)} \leq 1$ units of flow on each (of $d(e)$) copies of the arc (s_p, t_q) in \widehat{G} , is a feasible $\widehat{\sigma}$ -flow with cost $\frac{|b|_1}{2}$. (Recall that in the minimum-cost problem we assume that arc capacities are infinite, thus feasibility condition (2) boils down to non-negativity of all f_e^* s.) So, we can assume that our $\widehat{\sigma}$ -flow \mathbf{f} has its cost $\widehat{\ell}(\mathbf{f})$ at most $\frac{|b|_1}{2} + \frac{1}{2}$. (Otherwise, we know that there is no perfect \mathbf{b} -matching in G .)

Now, given any feasible $\widehat{\sigma}$ -flow in \widehat{G} , we can decompose it into a collection of flow-paths and flow-cycles, where each of these flow-paths transports some amount of flow from some vertex s_p to some vertex t_q . By our construction of the graph \widehat{G} , each such flow-path has to have a length at least 1. On the other hand, if this flow-path is indeed of length exactly 1 then it has to correspond to a single arc (s_p, t_q) that reflects the existence of edge (p, q) in G . As a result, our feasible $\widehat{\sigma}$ -flow \mathbf{f} in \widehat{G} has to have its cost $\widehat{\ell}(\mathbf{f})$ to be at least $\frac{|b|_1}{2}$ and, furthermore, $\widehat{\ell}(\mathbf{f}) - \frac{|b|_1}{2}$ is an upper bound on the total amount of flow in \mathbf{f} that is not transported over the direct one-arc flow paths (and thus passes through the vertex v^*).

So, as we argued that the cost of \mathbf{f} has to be at most $\frac{|b|_1}{2} + \frac{1}{2}$, there is only at most $\frac{1}{2}$ units of flow in \mathbf{f} that passes through the vertex v^* . Now, to extract the desired (fractional) near-perfect \mathbf{b} -matching \mathbf{x} , we just take $x_e = f_{(s_p, t_q)}$, for each edge $e = (p, q)$ in G . Clearly, the size of such fractional matching is at least $\frac{|b|_1}{2} - \frac{1}{2}$, which is well above our lowerbound of $\frac{|b|_1}{2} - \widetilde{O}(m^{\frac{3}{7}})$ for a near-perfect matching. Also, our construction works in $\widetilde{O}(\widehat{m})$ time, as desired. \square

Slack Variables

In the light of the above, our goal now is to compute the near-optimal solution to our minimum-cost $\widehat{\sigma}$ -flow problem instance in the graph \widehat{G} . Our approach to this task is inspired by so-called path-following interior-point methods [Ye97, Wri97, BV04]. At a very high level, we will start with certain initial solution that is far from being optimal, and then we will gradually improve – in an

iterative manner – its quality until close-to-optimal solution is obtained. This gradual improvement will be performed in a very specific way. It will always try to push the current solution further down so-called central path.

Before we can define the central path, let us first mention that, in general, there are two natural ways of tracking the progress of a current solution towards optimality. One of them is purely primal and relies on just maintaining a feasible solution \mathbf{f} and comparing its cost against some estimate of the cost of the optimal solution. The second one – and the one that we will actually use here – is based on primal-dual paradigm. Namely, in addition to maintaining a feasible primal solution \mathbf{f} , we will also keep a dual feasible solution \mathbf{y} . This dual solution provides an embedding of all the vertices in \widehat{G} into a line, i.e., \mathbf{y} just assigns a real number y_v to each vertex v of \widehat{G} . Its feasibility condition is that for any arc $e = (v, w)$ of \widehat{G} it should be the case that its *slack variable* $s_e := \hat{l}_e - y_w + y_v$ is always non-negative, i.e., that the length of the arc e in this embedding is never larger than its length according to the length vector $\hat{\mathbf{l}}$.

Before we proceed further, we note that the dual solution \mathbf{y} is uniquely determined – up to a translation – by the vector \mathbf{s} (given the length vector $\hat{\mathbf{l}}$). So, for notational convenience, from now on, we will describe the dual solutions in terms of the vector \mathbf{s} instead of \mathbf{y} .

Duality Gap

It is not hard to see that any feasible dual solution \mathbf{s} provides a lower-bound on the cost of the optimal solution (after all, this is just a consequence of weak duality). In particular, one has that for any pair (\mathbf{f}, \mathbf{s}) of feasible primal and dual solutions, the so-called *duality gap*, i.e., the difference between the upper bound on the value of optimal solution that is provided by the primal solution \mathbf{f} and the lower bound provided by the dual solution \mathbf{s} is exactly

$$\mathbf{f}^T \mathbf{s} = \boldsymbol{\mu}^T \mathbf{1} = \sum_e \mu_e,$$

where $\mu_e := f_e s_e$, for each arc e , and $\mathbf{1}$ is all-ones vector (of dimension \widehat{m}).

This means that one can obtain a close-to-optimal solution by devising a procedure that (quickly) converges to a pair of primal and dual solutions (\mathbf{f}, \mathbf{s}) whose duality gap $|\boldsymbol{\mu}|_1$ is small (in our case, at most $\frac{1}{2}$).

γ -Centered Solutions and the Central Path

To describe in more detail the convergence process we will employ, let us associate with each arc e a *measure* $\nu_e \geq 1$. One can view ν_e as a certain notion of importance of a given arc. (The motivation behind introducing this notion will be clear later.) We will always make sure that the measures of arcs are not smaller than 1 and also that their total sum is never too large. That is, we will make sure to maintain the following invariant.

Invariant 5.2. *We have that $\boldsymbol{\nu}^T \mathbf{1} = \sum_e \nu_e \leq 4\widehat{m}$ and for each arc e , $\nu_e \geq 1$.*

We want to note that when discussing the preservation of the above invariant we will only focus on ensuring that the upperbound is not violated. The fact that $\nu_e \geq 1$ for all arcs e will be automatically enforced as we will make sure that the initial measure of all the arcs is always at least 1 and our algorithm will never decrease any measures – they only might increase.

γ -centered solutions. Now, let us define, for each arc e , $\hat{\mu}_e := \frac{\mu_e}{\nu_e} = \frac{f_e s_e}{\nu_e}$ to be the normalized value of μ_e and let

$$\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) := \frac{\sum_e f_e s_e}{\sum_e \nu_e} = \frac{\sum_e \mu_e}{\sum_e \nu_e} = \frac{\sum_e \nu_e \hat{\mu}_e}{\sum_e \nu_e} \quad (10)$$

be the weighted average value of $\hat{\mu}_e$ with weights given by the measures $\boldsymbol{\nu}$.

We will call a solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ (where $\boldsymbol{\nu}$ represents the associated measures) γ -centered, for some $\gamma \geq 0$, if

$$\|\hat{\boldsymbol{\mu}} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})\mathbf{1}\|_{\boldsymbol{\nu}, 2} = \sqrt{\sum_e \nu_e (\hat{\mu}_e - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}))^2} \leq \gamma \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}), \quad (11)$$

where, for a given vector $\mathbf{x} \in \mathbf{R}^{\hat{m}}$,

$$\|\mathbf{x}\|_{\boldsymbol{\nu}, p} := \left(\sum_e \nu_e x_e^p \right)^{\frac{1}{p}}, \quad (12)$$

i.e., $\|\mathbf{x}\|_{\boldsymbol{\nu}, p}$ is the ℓ_p -norm of the vector \mathbf{x} reweighed by the measures $\boldsymbol{\nu}$.

Note that in a 0-centered solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ we have all $\hat{\mu}_e$ equal to $\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$. More generally, a simple but very useful observation is that

Fact 5.3. *For any γ -centered solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ we have that*

$$(1 - \gamma)\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \leq \hat{\mu}_e = \frac{f_e s_e}{\nu_e} \leq (1 + \gamma)\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}),$$

for each arc e .

$\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ as a measure of progress. The quantity $\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ will be important to us for one more reason. It will constitute our measure of progress on the quality of our maintained solution. To see why it indeed can serve this role, recall that by Invariant 5.2 we have that

$$\mathbf{f}^T \mathbf{s} = \sum_e \mu_e = \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \left(\sum_e \nu_e \right) \leq 4\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \hat{m}. \quad (13)$$

So, if our goal is to obtain a solution whose duality gap is at most $\frac{1}{2}$ we just need to make sure that the corresponding value of $\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ is at most $\frac{1}{8\hat{m}}$.

The main reason why we choose to measure our progress in terms of $\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ instead of the actual duality gap $\mathbf{f}^T \mathbf{s}$ is that in our algorithm we will sometime end up increasing measures of arcs. Such increases lead to an increase of the duality gap, so measuring our progress in terms of $\mathbf{f}^T \mathbf{s}$ would require dealing with such local non-monotonicity of this quantity. Continently, once we focus on keeping track of $\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ (and ensure that Invariant 5.2 is never violated), these issues will be avoided.

The central path. Finally, after introducing the above definitions, we can define the central path to be the set of all the 0-centered solutions.⁸ One can show that this set constitutes an actual path

⁸Strictly speaking, in the literature, the central path corresponds to 0-centered solutions with the measures of all arcs being one.

in feasible space that spans all the 0-centered solutions and passes arbitrarily close to (but never reaches) an optimal solution to our minimum cost flow problem. This explains the name of “path-following” interior-point methods, as they start with some initial 0-centered solution and gradually advance along the central path to get increasingly more optimal γ -centered solution for some small fixed γ .

Traversing the Central Path with Electrical Flows

Motivated by this path-following approach, our algorithm for computing near-optimal solution to the minimum-cost $\hat{\sigma}$ -flow problem will start with some 0-centered solution $(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0)$ that has fairly large value of $\hat{\mu}(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0)$ (and thus is far from being optimal). Then, we will devise a sequence of solutions $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$, where t is the step index, that have increasingly smaller value of $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ (and thus, indirectly, the duality gap) while making sure that they always are $\hat{\gamma}$ -centered for some small constant $\hat{\gamma} := \frac{1}{400}$. This way, our algorithm will eventually converge to the desired close-to-optimal solution.

To implement this approach, we start with the following lemma that shows we can get the initial 0-centered solution $(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0)$ – its proof appears in Appendix D.1.

Lemma 5.4. *There exists an explicit 0-centered primal-dual feasible solution $(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0)$ with $\sum_e \nu_e^0 \leq 3\hat{m}$ and $\hat{\mu}(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0) = 1$.*

Note that the bound on the total measure of the arcs ensures that the Invariant 5.2 is preserved. Furthermore, there is a slack of at least \hat{m} remaining between $\sum_e \nu_e^0$ and the upperbound of $4\hat{m}$ from Invariant 5.2. It will be used to accommodate future measure increases in our improved algorithm (see Section 6).

We now proceed to explaining how given some $\hat{\gamma}$ -centered solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$, we can modify it to obtain a $\hat{\gamma}$ -centered solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ that has a smaller value of $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$.

The associated flow $\hat{\mathbf{f}}^t$. For a given solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ let us call it σ -feasible, for some demand vector σ , if it is dual feasible (i.e., $\mathbf{s} \geq 0$) and if \mathbf{f} is a feasible σ -flow. (So, a $\hat{\sigma}$ -feasible solution is a solution that is primal-dual feasible for our minimum-cost $\hat{\sigma}$ -flow problem.) Next, given a σ -feasible solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$, let us define an *associated electrical flow* $\hat{\mathbf{f}}$ to be the electrical σ -flow in (the undirected projection of) \hat{G} determined by resistances \mathbf{r} that are given as

$$r_e := \frac{s_e}{f_e} = \frac{\mu_e}{(f_e)^2}, \quad (14)$$

for arc e . (Whenever we use this definition, it will be always the case that all f_e s are positive and thus the resistances r_e are well-defined.)

Making an improvement step. The central object in our procedure for taking an improvement step will be the electrical flow $\hat{\mathbf{f}}^t$ that is associated with the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$. The fundamental property of this flow is that it allows us to simultaneously update our solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ both in the primal (flow) space – via the flow $\hat{\mathbf{f}}^t$ itself – and in the dual (line embedding) space – via the vertex potentials $\hat{\phi}^t$ that induced $\hat{\mathbf{f}}^t$ (see (5)). (In Section 7, we provide a detailed description of the whole improvement step.)

As we will see, such a guided update not only decreases the duality gap of our solution, but also perfectly maintains its centering when only first-order terms (i.e., terms linear in the updates) are accounted for. Unfortunately, the second-order terms (i.e., the ones depending on the product of primal and dual updates) can disturb the centering. So, to be able to control this deficiency, we need to ensure that the step size δ^t that governs the “aggressiveness” of the improvement step is sufficiently small.

Of course, on the other hand, it is important for us to have this step be as large as possible. After all, the extent of our duality gap improvement – and thus overall convergence rate of our algorithm – is directly proportional to this size. So, it is crucial for us to develop a good grasp on how the size of that step relates to the properties of the flow $\hat{\mathbf{f}}^t$.

To this end, let us define, for some – not necessarily feasible – flow \mathbf{f} and a positive vector $\mathbf{x} > 0$, $\boldsymbol{\rho}(\mathbf{f}, \mathbf{x})$ to be the vector of congestions inflicted in \widehat{G} by \mathbf{f} with respect to capacities given by \mathbf{x} . That is,

$$\rho(\mathbf{f}, \mathbf{x})_e := \frac{|f_e|}{x_e}, \quad (15)$$

for each arc e in \widehat{G} .

Now, in Section 7, we present a precise implementation and analysis of our update step. (This implementation can be viewed as a direct analogue of the update steps of path-following interior-point methods.) The result of this analysis is presented in the following theorem, which, in particular, ties the congestion vector $\boldsymbol{\rho}(\hat{\mathbf{f}}^t, \mathbf{f}^t)$ inflicted by the electrical flow $\hat{\mathbf{f}}^t$ with respect to the primal solution \mathbf{f}^t , to an upperbound on the size δ^t of the improvement step.

Theorem 5.5. *Let $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ be a solution that is $\hat{\gamma}$ -centered and $\hat{\boldsymbol{\sigma}}$ -feasible, and let $\hat{\mathbf{f}}^t$ be the associated electrical flow. We can compute in $\tilde{O}(\widehat{m})$ time a $\hat{\gamma}$ -centered and $\hat{\boldsymbol{\sigma}}$ -feasible solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ with $\hat{\boldsymbol{\mu}}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \leq (1 - \delta^t)\hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$, as long as,*

$$0 < \delta^t \leq \min \left\{ \frac{\sqrt{\hat{\gamma}}}{\|\boldsymbol{\rho}(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\boldsymbol{\nu}^t, 4}}, \frac{1}{2} \right\}.$$

Furthermore, we have that the measures do not change, i.e., $\boldsymbol{\nu}^{t+1} = \boldsymbol{\nu}^t$, and if for each arc e , we define $(1 + \kappa_e^t) := \frac{(1-\delta^t)s_e^{t+1}f_e^t}{f_e^{t+1}s_e^t} = \frac{(1-\delta^t)r_e^{t+1}}{r_e^t}$ and $(1 + \bar{\kappa}_e^t) := \frac{(1-\delta^t)f_e^t}{f_e^{t+1}}$ to make $\boldsymbol{\kappa}^t$ (resp. $\bar{\boldsymbol{\kappa}}^t$) reflect the relative change (scaled by $(1 - \delta^t)$) of resistances \mathbf{r}^t (resp. flows \mathbf{f}^t) then $\|\boldsymbol{\kappa}^t\|_\infty, \|\bar{\boldsymbol{\kappa}}^t\|_\infty \leq \frac{1}{2}$ and

$$|\kappa_e^t|, |\bar{\kappa}_e^t| \leq 4(\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e + \hat{\kappa}_e^t),$$

for some vector $\hat{\boldsymbol{\kappa}}^t$ with $\|\hat{\boldsymbol{\kappa}}^t\|_{\boldsymbol{\nu}^t, 2} \leq \frac{1}{16}$.

So, we see that the allowed size δ^t of the improvement steps is proportional to how much the guiding flow $\hat{\mathbf{f}}^t$ resembles the current primal solution \mathbf{f}^t . Thus, for example, if there is some arc e that flows much larger flow in $\hat{\mathbf{f}}^t$ than in \mathbf{f}^t , i.e., an arc e with large value of $\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e$, this arc will be severely penalized by the ℓ_4 -norm measuring the quality of the resemblance.

Also, it is worth pointing out that it is very important that the above bound is based on ℓ_4 instead, say ℓ_2 norm. In fact, one can show (see Lemma 6.8) that in case of our problem the ℓ_2 norm of congestion vector is always $\Omega(\widehat{m}^{\frac{1}{2}})$. So, using ℓ_2 norm would not lead to any improvement over the $\Omega(\widehat{m}^{\frac{1}{2}})$ iteration bound.

5.1 Bounding the Running Time

At this point, we want to present a fairly elementary proof of $\delta := \Omega(\widehat{m}^{-\frac{1}{2}})$ lowerbound on our allowed improvement step size δ^t . Note that once we achieve that then, by Lemma 5.4 and Theorem 5.5, we will have that the value of our measure of progress $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ after T steps is at most

$$\hat{\mu}(\mathbf{f}^T, \mathbf{s}^T, \boldsymbol{\nu}^T) \leq \prod_{t=1}^T (1 - \delta^t) \leq (1 - \delta)^T.$$

So, by setting $T := \delta^{-1} \log 8\widehat{m} = \widetilde{O}(\widehat{m}^{\frac{1}{2}})$, we recover the $O(\widehat{m}^{\frac{1}{2}})$ iterations convergence bound of interior-point methods. This leads to a simple $\widetilde{O}(\widehat{m}\delta^{-1}) = \widetilde{O}(m^{\frac{3}{2}})$ -time procedure that produces a solution with duality gap at most

$$4\widehat{m}\hat{\mu}(\mathbf{f}^T, \mathbf{s}^T, \boldsymbol{\nu}^T) \leq 4\widehat{m}(1 - \delta)^T \leq \frac{1}{2},$$

where we used Invariant 5.2 (see (13)). This, in turn, by Lemma 5.1 provides us with a solution to our instance of near-perfect \mathbf{b} -matching problem.

Therefore, to conclude the analysis of the simple $\widetilde{O}(m^{\frac{3}{2}})$ -time algorithm for the near-perfect \mathbf{b} -matching problem, it remains to establish the claimed lowerbound on δ^t .

Congestion and energy. By Theorem 5.5, performing such lowerbounding of δ^t boils down to upperbounding $\|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\boldsymbol{\nu}^t, 4}$. To understand how the latter can be done, one should observe the following simple but crucial fact. (This fact follows from Fact 5.3 and definition of the resistances \mathbf{r}^t (14).)

Fact 5.6. *For any γ -centered solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and any flow $\hat{\mathbf{f}}$ in \widehat{G} we have that*

$$r_e^t \hat{f}_e^2 = \frac{s_e^t}{f_e^t} \hat{f}_e^2 \geq (1 - \gamma) \nu_e^t \frac{\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)}{(f_e^t)^2} \hat{f}_e^2 = (1 - \gamma) \nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \rho(\hat{\mathbf{f}}, \mathbf{f}^t)_e^2,$$

and, similarly,

$$r_e^t \hat{f}_e^2 \leq (1 + \gamma) \nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \rho(\hat{\mathbf{f}}, \mathbf{f}^t)_e^2,$$

for any arc e in \widehat{E} .

Observe that the above inequalities state that – up to a $(1 \pm \gamma)$ factor – the square of the congestion $\rho(\hat{\mathbf{f}}, \mathbf{f}^t)_e$ incurred by an arc e in the flow $\hat{\mathbf{f}}$ is upperbounded by

$$\frac{r_e^t \hat{f}_e^2}{\nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)},$$

which corresponds to normalized (by $\nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$) contribution of the arc e to the energy $\mathcal{E}_{\mathbf{r}^t}(\hat{\mathbf{f}})$ of the flow $\hat{\mathbf{f}}$ with respect to resistances \mathbf{r}^t .

This simple connection between the congestion of an arc in $\hat{\mathbf{f}}$ and its contribution to the energy of that flow that is provided by Fact 5.6 will be fundamental to the rest of our discussion. In particular, it gives us an intuition on why we even expect the guiding electrical flows $\hat{\mathbf{f}}^t$ to inflict small congestion with respect to \mathbf{f}^t and thus allow us to take a larger step size. This intuition is

based on an observation that the main goal of electrical flows is to minimize energy. So, by choosing the resistances appropriately, we in some sense align this goal with our goal of making as large step sizes as possible. Roughly speaking, we are employing here the ℓ_2 norm minimization offered by electrical flows to achieve the desired ℓ_4 -minimization corresponding to larger step sizes.

Now, an immediate consequence of the above connection is an elementary way of upperbounding $\|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\nu^t, 4}$: we just bound the ℓ_2 -energy of the guiding electrical flow $\hat{\mathbf{f}}^t$ that is associated with our solution $(\mathbf{f}^t, \mathbf{s}^t, \nu^t)$ and exploit the generic relationship between ℓ_2 and ℓ_4 norms.

To implement this approach, let us start with the following lemma that gives us a bound on the ℓ_2 -energy of the electrical flow $\hat{\mathbf{f}}^t$.

Lemma 5.7. *For any σ -feasible solution $(\mathbf{f}, \mathbf{s}, \nu)$ and associated electrical flow $\hat{\mathbf{f}}$, we have that*

$$\mathcal{E}_r(\hat{\mathbf{f}}) \leq \mathcal{E}_r(\mathbf{f}) \leq 4\hat{m}\hat{\mu}(\mathbf{f}, \mathbf{s}, \nu).$$

Proof. The fact that $\mathcal{E}_r(\hat{\mathbf{f}}) \leq \mathcal{E}_r(\mathbf{f})$ follows directly from the definition of $\hat{\mathbf{f}}$ and the fact that electrical σ -flow minimizes energy among all the σ -flows (which includes \mathbf{f}).

Now, by definition (14) of r and that of $\hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)$ (10) we have that

$$\mathcal{E}_r(\mathbf{f}) = \sum_e \frac{s_e}{f_e} f_e^2 = \sum_e s_e f_e = \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu) \left(\sum_e \nu_e \right) \leq 4\hat{m}\hat{\mu}(\mathbf{f}, \mathbf{s}, \nu),$$

where the last line follows by Invariant 5.2. □

Once we establish this upperbound on ℓ_2 -energy, we simply use it to upperbound the ℓ_4 -energy of the congestion vector. Specifically, by applying Cauchy-Schwarz inequality and the fact that $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_1$, for any vector \mathbf{x} , we get that

$$\|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\nu^t, 4}^4 = \sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^4 \leq \left(\sum_e \sqrt{\nu_e^t} \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^2 \right)^2 \leq \left(\sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^2 \right)^2 = \|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\nu^t, 2}^4,$$

where we also used the fact that $\nu_e^t \geq 1$.

Now, to bound the ℓ_2 norm (instead of ℓ_4 norm) of the congestion vector, we just note that by Fact 5.6 and Lemma 5.7

$$\left(\sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^2 \right)^2 \leq \left(\sum_e \frac{\nu_e^t r_e^t f_e^2}{(1-\hat{\gamma})\nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \nu^t)} \right)^2 = \left(\frac{\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)}{(1-\hat{\gamma})\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \nu^t)} \right)^2 \leq O(\hat{m}^2). \quad (16)$$

Therefore, we can conclude with the following lowerbound on δ^t .

Fact 5.8. *For any t , $\delta^t \geq \frac{1}{O(\sqrt{\hat{m}})}$.*

It is worth pointing out that, as we discussed before, the fact that we settled here for an ℓ_2 -norm-based (instead of an ℓ_4 -norm-based) dependence of δ^t on the congestion vector $\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)$, this $\frac{1}{O(\sqrt{\hat{m}})}$ lowerbound is the best possible to achieve with this approach. Therefore, to have any hope of obtaining an improvement that goes beyond this bound (as we will do in the next section), we crucially require to be working with ℓ_4 -norm-based (instead of only ℓ_2 -norm-based) arguments.

6 An Improved Algorithm for Bipartite b -Matching Problem

After setting up our primal-dual framework and presenting the $\tilde{O}(m^{\frac{3}{2}})$ -time algorithm in the previous section, we can now proceed to developing our improved algorithm with the running time of $\tilde{O}(m^{\frac{10}{7}}) = \tilde{O}(m^{\frac{3}{2}-\eta})$, for $\eta := \frac{1}{14} - o(1)$.

Given our analysis and discussion in the previous section, a tempting approach to obtaining such an improved bound could be trying to simply tighten our analysis performed there (e.g., by taking advantage of ℓ_4 -norm-based instead of only ℓ_2 -norm-based arguments) and thus improve the worst-case lowerbound on δ^t that we established (cf. Fact 5.8).

It turns out, however, that getting an improved bound is not merely a matter of performing a better analysis. In the worst-case, our $O(\hat{m}^{-\frac{1}{2}})$ bound is actually tight. After all, if there is an arc that incurs $\Omega(\hat{m}^{\frac{1}{2}})$ congestion in the associated electrical flow, the resulting ℓ_4 -norm of the congestion vector will be $\Omega(\hat{m}^{\frac{1}{2}})$. So, even though the connection between congestion and energy we established before (see Fact 5.6 and Lemma 5.7) tells us that there cannot be too many such arcs (as each one of them would need to contribute a very significant, $\Omega(1)$, fraction of the total energy of the electrical flow), having just one such arc is already enough to prevent us from taking larger than $O(\hat{m}^{-\frac{1}{2}})$ improvement step.

Therefore, as we cannot rule out that such worst-case situation arises in each iteration of our algorithm⁹, getting our desired improvement requires developing a strategy that explicitly ensures that this is not the case (or, at least, not too often).

At a high level, our general approach to accomplishing this goal is based on “massaging” the solution that we maintain. That is, we devise and carefully combine two methods of altering our solutions. These methods, on one hand, significantly improve the behavior of the associated electrical flow while, on other hand, only slightly perturb the characteristics of that solution that are vital to recovering the desired near-perfect b -matching at the end.

The first of these two methods is related to edge removal technique of Christiano et al. [CKM⁺11]. Their technique is based on repeated removal from the graph of the edges that suffer too much congestion. As [CKM⁺11] showed (via a simple energy-based argument), when such edge removal is applied to electrical flows that guide multiplicative-weight-update-based optimization routine, one obtains a significantly faster convergence to approximately optimal solution.

Unfortunately, as our primal-dual framework has much more delicate nature than the multiplicative-weight-update method, such removal of “bottlenecking” arcs would be too drastic and, in particular, could destroy the structure of our dual solution. Therefore, we apply a more careful approach.

Instead of removing arcs, we only perturb them by moderately increasing their lengths (and thus their slack variables). (Note that, by (14), increasing arc’s slack variable increases its resistance.) Furthermore, to avoid significant distortion of the dual solution, we do not apply this perturbation to all “bottlenecking” arcs, but only the ones that are “heavy” in the primal solution (see Definition 6.7 below).

We then use a certain refinement of the original energy-based argument of Christiano et al. [CKM⁺11] (that needs, in particular, to deal with the fact that – in contrast to the multiplicative-weight-update-based framework of [CKM⁺11] – in our framework the arc’s resistances can change

⁹One would suspect, however, that such situations are indeed rare. This might be one explanation of why in practice interior-point methods are able to take most of its step sizes to be very large and thus converge much faster than indicated by the worst-case bound of $O(\hat{m}^{\frac{3}{2}})$.

in a completely non-monotonic fashion) to show that the behavior of our guiding electrical flows on such “heavy” arcs is indeed improved.

Now, our second method – that is somewhat complementary to the first one and aims at accommodating the “light” arcs – is based on an appropriate preconditioning of our solution by augmenting it with auxiliary arcs. The purpose of adding these arcs is to improve the connectivity (and thus electrical conductance) of the underlying solution (when treated as a graph with resistances) while changing the structure of our original solution in only minimal way (that can be fixed later). We then show via a certain dual-based argument that existence of these auxiliary arcs ensures that “light” arcs are never the bottlenecking ones (and thus do need to be dealt with anymore).

We proceed now to detailed description and analysis of our improved algorithm.

The Sets $S_l(\hat{\mathbf{f}})$ and θ -Smoothness

We start by specifying the behavior of associated electrical flows that is “good” from our perspective. To this end, for a flow $\hat{\mathbf{f}}$ in \hat{G} , a solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$, and integer l , let us define $S_l(\hat{\mathbf{f}})$ to be the set of all the arcs e such that

$$\frac{\sqrt{\widehat{m}}}{2^{l+1}} < \rho(\hat{\mathbf{f}}, \mathbf{f})_e \leq \frac{\sqrt{\widehat{m}}}{2^l}, \quad (17)$$

i.e., the collection of all the arcs whose congestion in the flow $\hat{\mathbf{f}}$ (with respect to capacities given by \mathbf{f}) is between $\frac{\sqrt{\widehat{m}}}{2^{l+1}}$ and $\frac{\sqrt{\widehat{m}}}{2^l}$.

Now, we introduce a definition that will be fundamental to the rest of our discussion.

Definition 6.1. *For some $0 \leq \theta \leq 1$, a flow $\hat{\mathbf{f}}$, and solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ (that will be always clear from the context), we say that $\hat{\mathbf{f}}$ is θ -smooth on some of arcs $S \subseteq \widehat{E}$ iff, for any integer $l \leq \log \theta^{-3}$, we have that*

$$\boldsymbol{\nu}(S_l(\hat{\mathbf{f}}) \cap S) \leq \lfloor \theta^3 2^{3l} \rfloor,$$

where $\boldsymbol{\nu}(S') := \sum_{e \in S'} \nu_e$. Furthermore, we simply say that $\hat{\mathbf{f}}$ is θ -smooth if $S = \widehat{E}$, i.e., S contains all the arcs.

Clearly, the θ -smoothness constraints the distribution of the arcs that suffer high congestion in $\hat{\mathbf{f}}$. In particular, it implies that there is no arcs whose congestion $\rho(\hat{\mathbf{f}}, \mathbf{f})_e$ is larger than $\theta\sqrt{\widehat{m}}$.

Observe that the tight worst-case example for the lowerbound on δ^t (cf. Fact 5.8) corresponds to situation where the electrical flow $\hat{\mathbf{f}}^t$ associated with the maintained $\hat{\gamma}$ -centered solution solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ makes some arcs highly-congested, i.e., makes them suffer congestion of $\Omega(\sqrt{\widehat{m}})$. However, the above definition of θ -smoothness, forbids existence of such arcs. Therefore, the hope is that once our electrical flows $\hat{\mathbf{f}}^t$ are θ -smooth, a better lowerbound on δ^t (and thus faster convergence) is possible. As the following lemma shows, this hope is indeed well-founded.

Lemma 6.2. *Let $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ be a σ -feasible and $\hat{\gamma}$ -centered solution and let $\hat{\mathbf{f}}^t$ be the associated electrical flow that is θ -smooth, for some $0 \leq \theta \leq 1$. We have that*

$$\delta^t \geq \frac{1}{C_\delta \theta \sqrt{\widehat{m}}},$$

for some sufficiently large constant $C_\delta \geq 1$.

Proof. By Theorem 5.5, in order to lowerbound δ^t we need to upperbound the quantity

$$\|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\nu^t, 4}^4 = \sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^4.$$

To this end, note that

$$\sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^4 \leq \sum_l \nu^t(S_l(\hat{\mathbf{f}}^t)) \frac{\hat{m}^2}{2^{4l}} \leq \sum_l \lfloor \theta^3 2^{3l} \rfloor \frac{\hat{m}^2}{2^{4l}} \leq \sum_{l \geq \lfloor \log \theta^{-1} \rfloor} \frac{\theta^3 \hat{m}^2}{2^l} \leq 4\theta^4 \hat{m}^2,$$

where we used the θ -smoothness of $\hat{\mathbf{f}}^t$ (cf. Definition 6.1) and the fact that $\lfloor \theta^3 2^{3l} \rfloor = 0$ whenever $l < \lfloor \log \theta^{-1} \rfloor$. So, the lemma follows once $C_\delta > 0$ is chosen to be an appropriately large constant. \square

In the light of the above lemma, if we somehow knew that all – or, at least, most of – the flows $\hat{\mathbf{f}}^t$ that we compute are indeed θ -smooth for some $\theta = O(\hat{m}^{-\eta})$, we would immediately get the desired faster algorithm. Unfortunately, as we already discussed, it seems to be hard to argue that this is what happens in the worst-case. Therefore, to address this problem we develop a perturbation approach that we carefully apply to our maintained solutions to ensure that most of the flows $\hat{\mathbf{f}}^t$ is indeed θ -smooth for some small enough value of θ .

α -Stretching

One of the main operations that we will use to implement our perturbations is called α -stretching. To describe it, consider a solution $(\mathbf{f}, \mathbf{s}, \nu)$ that is γ -centered and a parameter $\alpha \geq 0$. We define an α -stretching of an arc e to be an operation that returns a solution $(\mathbf{f}', \mathbf{s}', \nu')$ obtained from $(\mathbf{f}, \mathbf{s}, \nu)$ by, first, increasing the length \hat{l}_e of the arc e (and thus the value of s_e) by αs_e and, then, increasing the measure ν_e of e by a factor of $(1 + \beta)$, where

$$\beta := \frac{\alpha f_e s_e}{\nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)}. \quad (18)$$

The remaining part of the solution remains the same.

The property of α -stretching that is key from our point of view, is that after applying it to some arc e its resistance $r_e := \frac{s_e}{f_e}$ increases by a factor of exactly $(1 + \alpha)$. Furthermore, our choice of value of β is justified by the lemma below – its proof appears in Appendix E.1.

Lemma 6.3. *If $(\mathbf{f}, \mathbf{s}, \nu)$ was a γ -centered solution with $\gamma \leq \frac{1}{2}$ then so will be $(\mathbf{f}', \mathbf{s}', \nu')$ and $\hat{\mu}(\mathbf{f}', \mathbf{s}', \nu') = \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)$. Furthermore, $(1 - \gamma)\alpha \leq \beta \leq (1 + \gamma)\alpha$.*

So, we see that applying α -stretching with this setting of β does not perturb our measure of progress $\hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)$, even though the duality gap $\mathbf{f}^T \mathbf{s}$ changes due to corresponding increase in measure. (Again, this is one reason why we chose $\hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)$ to measure our progress.)

On the other hand, besides the increase in measure, another undesirable byproduct of α -stretching is the increase of arc's length. To mitigate the effect of this process on the validity of our final solution, we will ensure that the following invariant is maintained throughout the algorithm.

Invariant 6.4. *The overall increase of arcs' length due to α -stretching is at most $\tilde{O}(\hat{m}^{\frac{1}{2}-\eta})$ and no individual arc has its total increase of length larger than 1.*

Maintaining this invariant will be important for two reasons. One is captured by the following simple lemma, whose proof appears in Appendix E.2.

Lemma 6.5. *If Invariant 6.4 is preserved then for any σ -feasible solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$, we have that s_e is at most 6, for any arc e .*

The other, and even more important one, is that as long as this invariant is preserved the final close-to-optimal solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ to our perturbed problem still allows us to recover the desired near-perfect \mathbf{b} -matching in the (original) graph G (or conclude that no perfect \mathbf{b} -matching exists). More precisely, in Appendix E.3 we prove the following lemma.

Lemma 6.6. *Provided Invariant 6.4 holds, given any feasible $\hat{\sigma}$ -flow \mathbf{f} in \hat{G} whose cost is within additive $\frac{1}{2}$ of the optimum, we can recover in $\tilde{O}(\hat{m})$ time a (fractional) near-perfect \mathbf{b} -matching in G , or conclude that no perfect \mathbf{b} -matching exists in G .*

Heavy Arcs

As we already mentioned, an important role in our improved algorithm is played by a classification of arcs into two classes, “heavy” and “light”, depending on their current flow in the primal solution. We make this classification precise below.

Definition 6.7. *Given a γ -centered solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ with $\gamma \leq \frac{1}{2}$, we call an arc e heavy if $f_e \geq \nu_e F_H$, where*

$$F_H := C_H^{-1} \hat{m}^{\frac{1}{2} - 3\eta} \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}),$$

for some sufficiently large constant $C_H > 1$ that we will fix later (see Lemma 6.15). We say that an arc is light if it is not heavy.

The motivation for the above classification stems from a desire to control the increase in arc’s length due to an application of α -stretching. Namely, observe that if we α -stretch an heavy arc e then the increase in this arc’s length is by at most

$$\alpha s_e \leq \frac{\alpha(1 + \gamma)\nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{f_e} \leq C_H(1 + \gamma)\alpha \frac{\hat{m}^{3\eta}}{\sqrt{\hat{m}}},$$

where we used Fact 5.3. So, as long as we apply α -stretching operations only to heavy arcs – which essentially will be the case in our algorithm – we can guarantee that the resulting change in arc length is relatively small. This will be important to ensuring that Invariant 6.4 is never violated.

Having introduced the above concepts, we are ready to proceed to presenting our improved algorithm. In this presentation, we fix for the rest of this section $\hat{\theta} := \hat{m}^{-\eta}$, where

$$\eta := \frac{1}{14} - C_\eta \frac{\log \log n}{\log n} \tag{19}$$

and C_η is a sufficiently large constant to be fixed later.

We describe our algorithm in two stages. First, in Section 6.1, we present a variant of the algorithm that works under an ad-hoc assumption that all the electrical flows $\hat{\mathbf{f}}^t$ that we compute are always $\hat{\theta}$ -smooth on the set of light arcs. (So, we need to deal there only with its possible non- $\hat{\theta}$ -smoothness on the set of heavy arcs.) Then, in Section 6.2, we show how to apply a preconditioning technique to obtain an augmented version of our graph such that when we run our algorithm it is indeed true that the above ad-hoc assumption holds.

6.1 Perturbing Heavy Arcs

In this section, we work under an ad-hoc assumption that the electrical flows $\hat{\mathbf{f}}^t$ that are associated with the maintained solutions $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ are always $\hat{\theta}$ -smooth – with $\hat{\theta} := \hat{m}^{-\eta}$ – on the set of arcs that are light with respect to that solution. We present an $\tilde{O}(\hat{m}^{\frac{3}{2}-\eta})$ -time algorithm for this setting.

$\hat{\theta}$ -Improvement Phase

The core of this algorithm is an implementation of a primitive we call a $\hat{\theta}$ -improvement phase. This primitive, given a $\hat{\boldsymbol{\sigma}}$ -feasible and $\hat{\gamma}$ -centered solution $(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$, returns in $\tilde{O}(\hat{m}^{1+2\eta})$ time a $\hat{\boldsymbol{\sigma}}$ -feasible and $\hat{\gamma}$ -centered solution $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ such that

$$\hat{\mu}(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f}) \leq \hat{\lambda} \hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0}), \quad (20)$$

where $\hat{\lambda} := \left(1 - \frac{1}{2C_\delta \hat{\theta} \sqrt{\hat{m}}}\right)^{\hat{\theta}^{-2}}$ and C_δ is the constant from Lemma 6.2.

Observe that once we obtain an implementation of such a $\hat{\theta}$ -improvement phase, we can get the desired improved algorithm as follows. We start with a $\hat{\boldsymbol{\sigma}}$ -feasible and $\hat{\gamma}$ -centered solution $(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0)$ as in Lemma 5.4. Next, we apply \hat{T} iterations of $\hat{\theta}$ -improvement phase to it, with

$$\hat{T} := 2C_\delta \hat{\theta}^3 \sqrt{\hat{m}} \ln 8\hat{m} = O(\hat{m}^{\frac{1}{2}-3\eta} \log \hat{m}). \quad (21)$$

Note that after doing this, we know that if $(\mathbf{f}^F, \mathbf{s}^F, \boldsymbol{\nu}^F)$ is the final $\hat{\boldsymbol{\sigma}}$ -feasible $\hat{\gamma}$ -centered solution we compute then

$$\hat{\mu}(\mathbf{f}^F, \mathbf{s}^F, \boldsymbol{\nu}^F) \leq \hat{\lambda}^{\hat{T}} \hat{\mu}(\mathbf{f}^0, \mathbf{s}^0, \boldsymbol{\nu}^0) = \left(1 - \frac{1}{2C_\delta \hat{\theta} \sqrt{\hat{m}}}\right)^{2C_\delta \hat{\theta}^3 \sqrt{\hat{m}} \ln 8\hat{m}} \leq \frac{1}{8\hat{m}}.$$

So, as long as we can show that $(\mathbf{f}^F, \mathbf{s}^F, \boldsymbol{\nu}^F)$ satisfies Invariants 5.2 and 6.4, we can use Lemma 6.6 to recover the desired near-perfect \mathbf{b} -matching in G or conclude that no perfect \mathbf{b} -matching exists in G . Also, the overall running time of this algorithm will indeed be $\tilde{O}(\hat{T} \hat{m}^{1+2\eta}) = \tilde{O}(\hat{m}^{\frac{3}{2}-\eta})$, as desired.

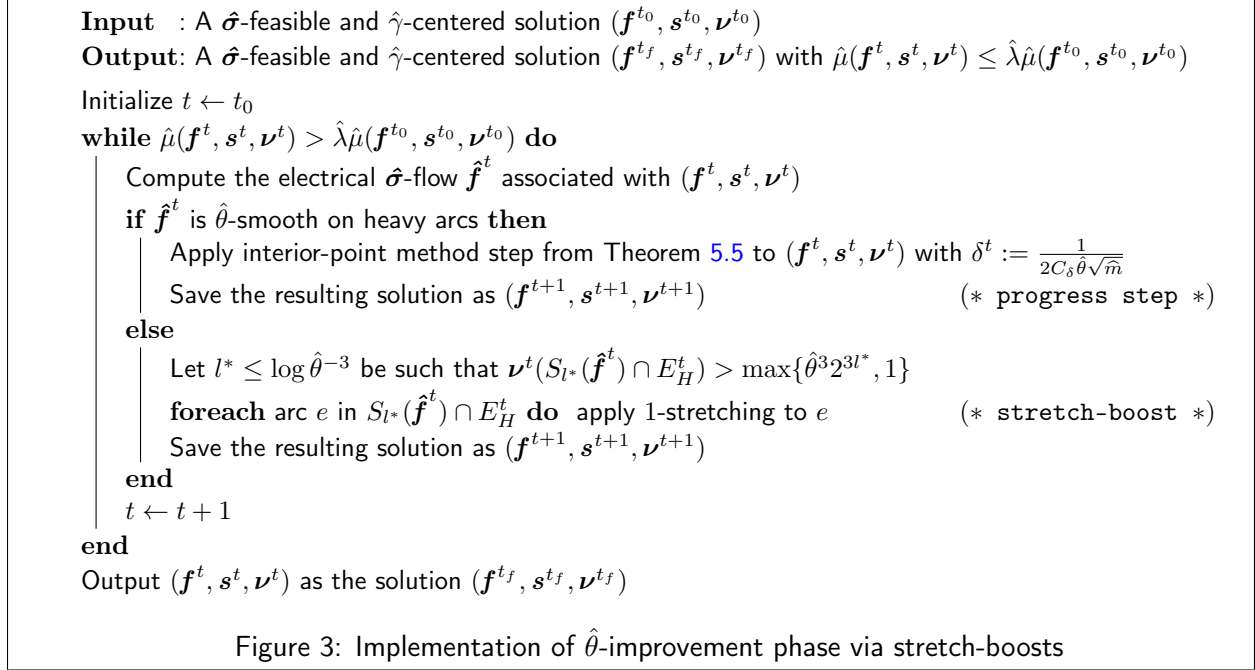
Implementation of $\hat{\theta}$ -Improvement Phase via Stretch-boosts

In the light of the above discussion, we just need to focus on implementing the $\hat{\theta}$ -improvement phase, as well as, ensuring that running it for \hat{T} iterations will not violate Invariants 5.2 and 6.4.

Our implementation – presented in Figure 3 – is an iterative procedure. We maintain a $\hat{\boldsymbol{\sigma}}$ -feasible $\hat{\gamma}$ -centered solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ – initially, $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ is equal to $(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$. Next, as long as $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) > \hat{\lambda} \hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ we repeat the following iterative step.

We first check if the electrical flow $\hat{\mathbf{f}}^t$ associated with $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ is $\hat{\theta}$ -smooth on the set of heavy arcs.

If it is indeed the case then one can easily see that such $\hat{\mathbf{f}}^t$ needs to be $2\hat{\theta}$ -smooth (on the set of all the arcs). (This uses our ad-hoc assumption that all $\hat{\mathbf{f}}^t$ we compute are always $\hat{\theta}$ -smooth on the set of light arcs.) So, in this situation, we can just apply an interior-point method step – as described in Theorem 5.5 – to $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ with setting $\delta^t := \frac{1}{2C_\delta \hat{\theta} \sqrt{\hat{m}}}$. (Note that by Lemma 6.2 this setting of δ^t is valid.) For future reference, we call this step a *progress step*. After executing it, we



set the resulting $\hat{\sigma}$ -feasible and $\hat{\gamma}$ -centered solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ as our current solution and proceed to next iterative step.

Otherwise, that is, if $\hat{\mathbf{f}}^t$ is not $\hat{\theta}$ -smooth on the set of heavy arcs, then – by Definition 6.1 – there is an $l^* \leq \log \hat{\theta}^{-3}$ such that

$$\boldsymbol{\nu}^t(S_{l^*}(\hat{\mathbf{f}}^t) \cap E_H^t) > \max\{\hat{\theta}^3 2^{3l^*}, 1\}, \quad (22)$$

where we used the fact that all measures of arcs are always at least one and E_H^t denotes the set of heavy arcs with respect to the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$.

To cope with this situation, we perform 1-stretching of all the arcs in $S_{l^*}(\hat{\mathbf{f}}^t) \cap E_H^t$. (Note that, by Lemma 6.3, this operation does not change the value of $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and our solution remains $\hat{\sigma}$ -feasible and $\hat{\gamma}$ -centered.) We call this operation *stretch-boosting* and l^* will be referred to as the *index* of this stretch-boosting. After performing stretch-boosting, we proceed to the next iterative step.

This finishes the description of our implementation.

Analysis

To analyze the above procedure, let us note that due to our stopping condition, once this procedure terminates the resulting solution $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ satisfies our requirements. Also, there will be at most $\hat{\theta}^{-2}$ progress steps executed. This is so, as 1-stretching does not affect the value of $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and, by Theorem 5.5, each progress step decreases $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ by a factor of at least $(1 - \delta^t) = \hat{\lambda} \hat{\theta}^{-2}$. Thus, as each of these steps runs in $\tilde{O}(\hat{m})$ time, the resulting total time of progress steps is $\tilde{O}(\hat{m} \hat{\theta}^{-2}) = \tilde{O}(\hat{m}^{1+2\eta})$, as desired.

Therefore, we can just focus on bounding the number of stretch-boost operations executed, as

well as, on showing that calling our implementation of $\hat{\theta}$ -improvement phase \hat{T} times does not violate Invariants 5.2 and 6.4.

We start with the former task. In this bounding of the number of stretch-boost operations, we assume that Invariant 5.2 holds. We will justify this assumption later when proving that our two desired invariants are indeed preserved by our algorithm.

To do the bounding, we consider the energy $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ of the electrical $\hat{\sigma}$ -flow $\hat{\mathbf{f}}^t$ (determined by resistances \mathbf{r}^t given by (14)) that is associated with our current solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$. We treat this quantity as a potential function and show the following facts:

- (a) $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ is always at least $C_E^{-1}\hat{m}\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and at most $C_E\hat{m}\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$, for some sufficiently large constant $C_E > 1$ – see Lemma 6.8;
- (b) $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ increases by a factor of at least $(1 + C_S\hat{\theta}^2)$, for some constant $C_S > 0$, whenever a stretch-boosting step is applied – see Lemma 6.9;
- (c) $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ decreases by a factor of at most $(1 + C_P\hat{\theta}^2 \ln \hat{m})$ each time a progress step is executed, where $C_P > 0$ is some sufficiently large constant – see Lemma 6.11.

Note that once the above statements are established, it must be the case that there is at most $T_s := C_P C_S^{-1} \hat{\theta}^{-2} \ln C_E^2 \hat{m} = \tilde{O}(\hat{\theta}^{-2})$ stretch-boost operation overall. To see that, assume this was not the case, i.e., that there was more than T_s stretch-boosts. Then, by the above statements and the fact that there is at most $\hat{\theta}^{-2}$ progress steps we would have that

$$\mathcal{E}_{r^{t_f}}(\hat{\mathbf{f}}^{t_f}) > \frac{(1 + C_S\hat{\theta}^2)^{T_s} \mathcal{E}_{r^{t_0}}(\hat{\mathbf{f}}^{t_0})}{(1 + C_P\hat{\theta}^2 \ln \hat{m})^{\hat{\theta}^{-2}}} \geq \frac{(1 + C_S\hat{\theta}^2)^{T_s} \hat{m}\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})}{C_E \hat{m}^{C_P}} \geq C_E \hat{m}\hat{\mu}(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f}),$$

which would violate the upperbound on energy $\mathcal{E}_{r^{t_f}}(\hat{\mathbf{f}}^{t_f})$ established by statement (a).

So, it must be then indeed the case that there is at most $T_s = \tilde{O}(\hat{\theta}^{-2})$ stretch-boosts, which gives the desired $\tilde{O}(\hat{\theta}^{-2}\hat{m}) = \tilde{O}(\hat{m}^{1+2n})$ total running time bound.

In the light of the above, we can turn our attention to proving statements (a)-(c). We start with statement (a). This statement essentially follow from Lemma 5.7 and some simple energy-lowerbounding argument. The prove of the following lemma appears in Appendix E.4.

Lemma 6.8. *Let $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ be a $\hat{\sigma}$ -feasible and $\hat{\gamma}$ -centered solution. Provided that Invariant 5.2 holds, we have that*

$$C_E^{-1}\hat{m}\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \leq \mathcal{E}_{r^t}(\hat{\mathbf{f}}^t) \leq C_E\hat{m}\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t),$$

where $\hat{\mathbf{f}}^t$ is the electrical $\hat{\sigma}$ -flow associated with the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and $C_E > 1$ is a sufficiently large constant.

Next, we proceed to analyzing the effect of stretch-boosting on the energy $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$. Intuitively, by the connection between congestion and energy hinted by Fact 5.6, we know that the arcs with large congestion have to have unusually high contribution to the energy $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$. So, as 1-stretching effectively doubles the resistances of such arcs, it is not surprising that it ends up significantly increasing that energy. We make this formal – and thus establish statement (b) – in the lemma below. Its proof appears in Appendix E.5.

Lemma 6.9. *Each stretch-boost increases $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ by a factor of at least*

$$\left(1 + C_S \hat{\theta}^2 \left(\nu^t(S_{l^*}(\hat{\mathbf{f}}^t) \cap E_H^t)\right)^{\frac{1}{3}}\right) \geq \left(1 + C_S \hat{\theta}^2\right),$$

for some constant $C_S > 0$.

To complete our analysis, it remains to show that our potential $\mathcal{E}_{r^t}(\hat{\mathbf{f}}^t)$ does not decrease too much during the progress steps. In other words, we prove statement (c).

Note that the difficulty here stems from the fact that, in principle, the resistances of arcs can change pretty arbitrarily during a progress step. They can either increase or decrease and even by a constant multiplicative factor, thus possibly leading to severe and very non-monotone energy fluctuations.

The key reason that enables us to control that energy change after all, is that we perform the progress step only if the flow $\hat{\mathbf{f}}^t$ is $\hat{\theta}$ -smooth. This is helpful in two ways. Firstly, because we can use it together with the connection between the change of the resistance of an arc and its congestion that we established in Theorem 5.5, to show that there is not too many arcs that significantly change their resistance (see Lemma 6.10 below). Secondly (and even more importantly), our connection between congestion and energy, allows us to conclude that $\hat{\theta}$ -smoothness implies that there is no small (measure-wise) set of arcs that contributes unusually high portion of the energy. So, even if some small set of arcs changes its resistances significantly, it is not able to influence the overall energy by too much (see Lemma 6.11). (In a sense, this intuition is one of the main motivations for introducing the notion of $\hat{\theta}$ -smoothness.) We, again, formalize this intuition below.

First, for a given vector λ and an integer l , let us define T_l^λ to be the set of all arcs such that

$$\frac{1}{2^{l+1}} \leq |\lambda_e| \leq \frac{1}{2^l}. \quad (23)$$

Now, we say that λ is τ -restricted, for some measure ν and $\tau \geq 0$ if for any $l \geq 0$,

$$\nu(T_l^\lambda) \leq \tau 2^{3l}. \quad (24)$$

Now, the lemma below bounds the change of resistances during any of our progress steps.

Lemma 6.10. *Let $\hat{\mathbf{f}}^t$ be a $2\hat{\theta}$ -smooth electrical flow associated with $(\mathbf{f}^t, \mathbf{s}^t, \nu^t)$. Let $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \nu^{t+1})$ be the solution obtained by applying an interior-point method step – as in Theorem 5.5 – to $(\mathbf{f}^t, \mathbf{s}^t, \nu^t)$ with $\delta^t := (2C_\delta \hat{\theta} \sqrt{\widehat{m}})^{-1}$. Then the vectors κ^t and $\bar{\kappa}^t$ are all C_R -restricted (with respect to ν^t) for some constant $C_R > 0$.*

Proof. We will prove that both the vector $\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)$ and the vector $\hat{\kappa}^t$ are $O(1)$ -restricted. It is easy to see that then the bound from Theorem 5.5 will imply that $\bar{\kappa}^t$ and κ^t are $O(1)$ -restricted too. So, choosing large enough constant C_R will prove the lemma.

To this end, observe $\hat{\kappa}^t$ is $O(1)$ -restricted as $\|\hat{\kappa}^t\|_{\nu^t, 2}$ by Theorem 5.5. On the other hand, note that if

$$\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e \geq \frac{1}{2^{l'}},$$

for some $l' \geq 1$ and arc e , then $e \in S_{l'}(\hat{\mathbf{f}}^t)$ for some

$$l' \leq l + \log \delta^t \sqrt{\widehat{m}} + 2 \leq l + \log \hat{\theta}^{-1} + O(1).$$

But by $2\hat{\theta}$ -smoothness of $\hat{\mathbf{f}}^t$, this means that the total measure of such arcs is at most

$$\lfloor \hat{\theta}^{-3} 2^{3l'} \rfloor \leq O(2^{3l}),$$

which establishes that $\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)$ is indeed $O(1)$ -restricted. The lemma follows. \square

Using the above observation, we can now finish establishing property (c) by proving the following lemma whose appears in Appendix E.6.

Lemma 6.11. *Let $\hat{\mathbf{f}}^t$ be a $2\hat{\theta}$ -smooth electrical flow associated with the $\hat{\gamma}$ -centered solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$. Let $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ be the solution obtained by applying an interior-point methods step – as in Theorem 5.5 – to $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ with $\delta^t := (2C_\delta \hat{\theta} \sqrt{\hat{m}})^{-1}$. Then,*

$$\mathcal{E}_{r^{t+1}}(\hat{\mathbf{f}}^{t+1}) \geq (1 + C_P \theta^2 \ln \hat{m})^{-1} \mathcal{E}_{r^t}(\hat{\mathbf{f}}^t),$$

where $\hat{\mathbf{f}}^{t+1}$ is the electrical flow associated with $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ and $C_P > 1$ is a sufficiently large constant.

Preservation of Invariants 5.2 and 6.4

Now, as we completed the analysis of the running time of our $\hat{\theta}$ -improvement phase implementation, we establish the remaining claim, i.e., we prove that executing the above procedure \hat{T} times does not lead to violation of Invariants 5.2 and 6.4.

Bounding measure increase. To this end, let us first focus on bounding the measure increases. By Lemma 6.3, we know that whenever we 1-stretch an arc e , its measure increases by at most $(1 + \hat{\gamma})\nu_e^t$. So, to bound the total measure increase it suffices to bound the total measure of arcs that are affected by 1-stretches across all the stretch-boost operations. (Here, if the same arc becomes 1-stretched multiple times, in different stretch-boosts, we account for its measure multiple times.)

In order to do that, note that by Lemma 6.9, if ν_i is the measure of the set of arcs that are 1-stretched in i -th stretch-boost, we have that the total increase of energy resulting from that is at least

$$\prod_{i=1}^k (1 + C_S \hat{\theta}^2 \nu_i^{\frac{1}{3}}).$$

Also, by Lemma 6.8, we know that we have to have that in any single stretch-boost, the energy cannot increase by more than C_E^2 factor. So, we have that

$$(1 + C_S \hat{\theta}^2 \nu_i^{\frac{1}{3}}) \leq C_E^2$$

and thus $\nu_i \leq \nu_{\max} := C_E^6 \hat{\theta}^6 = O(\hat{m}^{6\eta})$, for each i .

As a result, we can lowerbound the total increase of energy due to stretch-boosts by

$$\prod_{i=1}^k (1 + C_S \hat{\theta}^2 \nu_i^{\frac{1}{3}}) \geq (1 + C_S \hat{\theta}^2 \nu_{\max}^{\frac{1}{3}})^{\frac{\nu}{\nu_{\max}}},$$

where $\nu := \sum_i \nu_i$.

Finally, by Lemma 6.8 and Lemma 6.11, as well as, the fact that we have at most $\hat{\theta}^{-2}$ progress steps, we know that the overall (multiplicative) increase of energy resulting from all the stretch-boosts can be at most

$$(1 + C_S \hat{\theta}^2 \nu_{\max}^{\frac{1}{3}})^{\frac{\nu}{\nu_{\max}}} \leq \prod_{i=1}^k (1 + C_S \hat{\theta}^2 \nu_i^{\frac{1}{3}}) \leq C_E^2 (1 + C_P \hat{\theta}^2 \log \hat{m})^{\hat{\theta}^{-2}}.$$

Therefore, as $\hat{\theta}^2 \nu_{\max}^{\frac{1}{3}}$ is $\Omega(1)$, we have that the total measure increase ν is at most

$$\nu \leq \tilde{O}(\nu_{\max}) = \tilde{O}(\hat{m}^{6\eta}). \quad (25)$$

As a result, after executing at most \hat{T} $\hat{\theta}$ -improvement phases, the overall increase of measure can be bounded by

$$\hat{T} \cdot \tilde{O}(\hat{m}^{6\eta}) = \tilde{O}(\hat{m}^{\frac{1}{2}+3\eta}) < \hat{m}. \quad (26)$$

Now, given that by Lemma 5.4, we start with our measure being at most $3\hat{m}$ and thus have a slack of at least \hat{m} measure left before Invariant 5.2 becomes violated, this overall increase will indeed not lead to violation of this invariant.

Bounding arc length increase. To show that Invariant 6.4 is preserved as well, let us first note that the only way for length of arcs to increase is due to 1-stretching occurring during stretch-boosts. Furthermore, we only 1-stretch an arc if it is heavy. So, if a given (heavy) arc e gets 1-stretched at some step t then its length increases by at most

$$s_e^t \leq \frac{(1 + \hat{\gamma}) \nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)}{f_e^t} \leq \frac{(1 + \hat{\gamma}) C_H \hat{m}^{3\eta}}{\sqrt{\hat{m}}}.$$

On the other hand, by Lemma 6.3, the increase of measure of such arc is at least $(1 - \hat{\gamma}) \nu_e^t$. So, as $\nu_e^t \geq 1$, the increase of measure of an arc is within a factor of $O(\frac{\hat{m}^{3\eta}}{\sqrt{\hat{m}}}) = O(\hat{m}^{-4\eta})$ of increase of the length. So, as we just proved that the total measure increase is at most $\tilde{O}(\hat{m}^{\frac{1}{2}+3\eta})$ (cf. (26)), the desired bound of $\tilde{O}(\hat{m}^{\frac{1}{2}-\eta})$ on the total length increase follows.

Finally, as each 1-stretch increases the measure by a factor of at least $(2 - \hat{\gamma}) \geq \frac{3}{2}$ and – as we discussed above – we never 1-stretch anymore an arc whose measure is bigger than $\nu_{\max} = O(\hat{m}^{6\eta})$, no single arc will get 1-stretched more than $O(\log \nu_{\max}) = O(\log \hat{m})$ times. As a result, no single arc has its length increased by more than $O(\hat{m}^{-4\eta} \log \hat{m})$ that is much smaller than 1. Therefore, the Invariant 6.4 is also preserved. This concludes our analysis.

6.2 Preconditioning the Graph \hat{G}

Our analysis from the previous section was crucially relying on the assumption that all the flows $\hat{\mathbf{f}}^t$ are always $\hat{\theta}$ -smooth on the set of light arcs. Unfortunately, this assumption is not always valid.

To cope with this problem, we develop a modification of our algorithm that ensures that this $\hat{\theta}$ -smoothness assumption holds after all. Roughly speaking, we achieve that by an appropriate preconditioning our solution at the beginning of each $\hat{\theta}$ -improvement phase. This preconditioning is based on augmenting the graph \hat{G} with additional, auxiliary arcs and correspondingly extending our solution on them. These arcs are very light (i.e., have small value f_e^t of flow flowing through

them in augmented solution), while providing good connectivity (and thus relatively low effective resistance) between different vertices of the augmented graph.

The underlying intuition here is that the over-congestion of a light arc e is caused by amounts of flow that are at most $\sqrt{\widehat{m}}F_H$ (cf. Definition (6.7)) and thus are relatively small compared to the whole duality gap. So, by deploying these very light auxiliary arcs we encourage the electrical flow $\widehat{\mathbf{f}}^t$ to reroute such over-congesting flow from e and send it along auxiliary arcs. On the other hand, as the small value of this rerouted flow is small, the perturbation of our desired (non-augmented) solution introduced by these rerouting is relatively minor. Thus, we are able to deal with it relatively easily at the end of the whole $\widehat{\theta}$ -improvement phase, while still ending up making overall progress on the quality of our solution.

Augmenting the Graph and the Solution

The exact implementation of our preconditioning is based on modifying the execution of $\widehat{\theta}$ -improvement phase that was presented in the previous section in the following way. Let $(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ be the $\widehat{\gamma}$ -centered and $\widehat{\boldsymbol{\sigma}}$ -feasible solution at the beginning of some $\widehat{\theta}$ -improvement phase.

We start with augmenting the graph \widehat{G} by adding to it a new vertex \bar{v} , as well as, a_v copies of an arc (v, \bar{v}) and a_v copies of an arc (\bar{v}, v) , for each vertex v of \widehat{G} other than v^* , where

$$a_v := \sum_{e \in E(v)} \nu_e^{t_0} \quad (27)$$

is the total measure (with respect to $\boldsymbol{\nu}^{t_0}$) of all the arcs adjacent to v in \widehat{G} . We will call these newly added arcs *auxiliary* and denote the augmented graph as \widetilde{G} .

Next, we extend the solution $(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ to that augmented graph \widetilde{G} by assigning $f_e^{t_0} := F_A$, $s_e^{t_0} := \frac{\widehat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})}{f_e}$ and $\nu_e^{t_0} := 1$ to each auxiliary arc e , where

$$F_A := \frac{C_\delta \widehat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0}) \widehat{m}^{\frac{1}{2}-3\eta}}{C_A \widehat{T}} \quad (28)$$

with C_A being some sufficiently large constant to be fixed later, and the lengths of the auxiliary arcs being chosen so that the extended solution is still dual feasible. (As we will soon see, the actual lengths of auxiliary arcs are irrelevant.) Note that after this extension, the solution $(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ remains $\widehat{\gamma}$ -centered, $\widehat{\boldsymbol{\sigma}}$ -feasible and the value of $\widehat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ is unchanged. Also, observe that by Invariant 5.2, the number \bar{m} of arcs of the augmented graph \widetilde{G} is still only $O(\widehat{m})$. So, relating various quantities – in particular, the running times of our procedures – to either \widehat{m} or \bar{m} results in only a constant-factor discrepancy (that we will ignore in what follows).

Now, after the above preprocessing, we run the $\widehat{\theta}$ -improvement phase implementation, as described in the previous section, on the extended solution in the augmented graph \widetilde{G} . (In Section 6.2, we will prove that the assumption that underlies the analysis from the previous section, i.e., that all the flows $\widehat{\mathbf{f}}^t$ are $\widehat{\theta}$ -smooth on light arcs, is indeed valid.) The only further modification here is that after each progress step we $\widehat{m}^{2\eta} |\bar{\kappa}_e^t|$ -stretch each auxiliary arc e with $|\bar{\kappa}_e^t| \geq \widehat{\theta}^2$ (cf. Theorem 5.5). We will call this stretch operation *freezing*. (Note that as α -stretching only increases the resistances of arcs, this modification is compatible with the energy-based potential argument we employed in the previous section.) This freezing ensures that the flows on auxiliary arcs do not change to significantly in our solution and thus the impact of preconditioning provided by auxiliary

arcs on the quality of the final solution is minimized. We make this more precise in the following lemma whose proof appears in Appendix E.7.

Lemma 6.12. *During the whole $\hat{\theta}$ -improvement phase, we have that for each auxiliary arc e , $C_F^{-1}F_A \leq f_e^t \leq C_FF_A$, for some constant $C_F > 0$. Also, the total increase of measure of auxiliary arcs in that phase is at most $\tilde{O}(\hat{m}^{8\eta})$.*

Finally, once the execution of the above $\hat{\theta}$ -improvement phase finishes, we end up with a $\hat{\gamma}$ -centered and $\hat{\sigma}$ -feasible solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ such that $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \leq \hat{\lambda}\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$, as desired. However, this solution corresponds to the augmented graph \bar{G} instead of to the original graph \hat{G} .

To deal with this deficiency, we first simply discard all the auxiliary arcs and correspondingly truncate the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ to non-auxiliary arcs. Unfortunately, doing that might, in particular, render that solution not $\hat{\sigma}$ -feasible. So, to alleviate this problem, in Section 6.2 below, we describe a fixing procedure that, given such a truncated solution, produces the intended solution $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ that corresponds to the original graph \hat{G} , is $\hat{\gamma}$ -centered, $\hat{\sigma}$ -feasible and

$$\hat{\mu}(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f}) \leq \hat{\lambda}(1 + O(\hat{m}^{-\frac{1}{2}}))\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0}).$$

(Note that in our algorithm we are executing only $\hat{T} = \tilde{O}(\hat{m}^{\frac{1}{2}-3\eta})$ $\hat{\theta}$ -improvement phases overall. So, this additional $(1 + O(\hat{m}^{-\frac{1}{2}}))$ factor above is inconsequential.)

As we will see, a byproduct of this fixing procedure is an increase in the measure of (non-auxiliary) arcs. However, we will show that this increase is bounded by $O(\frac{C_FF\hat{m}}{C_A\hat{T}})$. Thus, taking C_A to be sufficiently large ensures that the resulting measure increases do not lead to violation of Invariant 5.2. (Note that the auxiliary arcs are always discarded at the end, so from the point of view of Invariant 5.2, it suffices that by Lemma 6.12 the measure of these arcs is always $o(\hat{m})$.)

In the light of the above discussion, all that remains is to describe and analyze the fixing procedure and to show that one can indeed assume that all the electrical flows $\hat{\mathbf{f}}^t$ computed during such $\hat{\theta}$ -improvement phase are $\hat{\theta}$ -smooth on the set of light arcs.

Fixing Procedure

We start by describing and analyzing the fixing procedure that we employ at the end of each $\hat{\theta}$ -improvement phase. Recall that in this procedure we are given as input a $\hat{\gamma}$ -centered and $\hat{\sigma}$ -feasible solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ in the augmented graph \bar{G} such that $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \leq \hat{\lambda}\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$. Our goal is to obtain a $\hat{\gamma}$ -centered $\hat{\sigma}$ -feasible solution $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ in the original graph \hat{G} that satisfies $\hat{\mu}(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f}) \leq \hat{\lambda}(1 + O(\hat{m}^{-\frac{1}{2}}))\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$.

We do this in two steps. First, we simply truncate the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ to the original graph \hat{G} by discarding all the auxiliary arcs and flow on them. Let us denote the resulting solution as $(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$. It is not hard to see that this solution is still $\hat{\gamma}$ -centered. In the following lemma – whose proof appears in Appendix E.8 – we argue that also the value of $\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$ has not increased by much.

Lemma 6.13. $\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') \leq \hat{\lambda}(1 + O(\hat{m}^{-\frac{1}{2}}))\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$.

At this point, we know that the solution $(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$ is $\hat{\gamma}$ -centered and $\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$ is as small as needed. Unfortunately, this solution can still be not $\hat{\sigma}$ -feasible.

Therefore, in the second step of our procedure, we address this last shortcoming. Our approach here requires introducing a certain simple operation. For a given some solution $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$, as well as,

some $\alpha \geq 0$ and an arc e , let us define α -widening of e (in $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$) as an operation in which we increase the value of f_e by a factor of $(1 + \alpha)$ and increase ν_e by a factor of $(1 + \beta)$, where β is given via (18).

We can view the α -widening operation as a counterpart of the α -stretching operation. In fact, one can see that due to symmetric nature of f_e and s_e and our choice of β , Lemma 6.3 also holds for α -widening operation. (Note that in the proof of Lemma 6.3 the roles of f_e and s_e are completely interchangeable.)

Now, our way of obtaining the desired solution $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ is very simple. Let us denote by $\bar{\boldsymbol{\sigma}}$ the actual demand vector of \mathbf{f}' and let $\tilde{\boldsymbol{\sigma}} := \hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}$ be the vector of demand differences. We start with $(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$ and for each vertex v of \hat{G} other than v^* , we do the following. If $\tilde{\sigma}_v \geq 0$ (resp. $\tilde{\sigma}_v < 0$), we apply α_v -widening to the arc $e(v) := (v, v^*)$ (resp. $e(v) := (v^*, v)$) with $\alpha_v := \frac{|\tilde{\sigma}_v|}{f'_e}$.

We take $(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f})$ to be the resulting solution. It is easy to see that this solution is $\hat{\boldsymbol{\sigma}}$ -feasible now. Also, by Lemma 6.3, we know that this solution remains $\hat{\gamma}$ -centered and that $\hat{\mu}(\mathbf{f}^{t_f}, \mathbf{s}^{t_f}, \boldsymbol{\nu}^{t_f}) = \hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') \leq \hat{\lambda}(1 + O(\hat{m}^{-\frac{1}{2}}))\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$, as needed.

So, we just need to establish the claimed bound of $O(\frac{C_F \hat{m}}{C_A \hat{T}})$ on total measure increase resulting from this procedure. To this end, note that by Lemma 6.3 this increase is at most

$$\begin{aligned} (1 + \hat{\gamma}) \sum_{v \neq v^*} \alpha_v \nu_{e(v)}^t &= (1 + \hat{\gamma}) \sum_{v \neq v^*} |\tilde{\sigma}_v| \frac{\nu_{e(v)}^t}{f'_{e(v)}} \leq (1 + \hat{\gamma})^2 \sum_{v \neq v^*} |\tilde{\sigma}_v| \frac{s'_{e(v)}}{\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')} \\ &= O\left(\sum_{v \neq v^*} \frac{|\tilde{\sigma}_v|}{\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')}\right) = O\left(\frac{|\hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}|_1}{\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')}\right), \end{aligned}$$

where we used the Fact 5.3 and Invariant 6.4, as well as, we applied Lemma 6.5 to conclude that each $s'_{e(v)}$ is $O(1)$.

Thus, in the light of the above, it only remains to bound $|\hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}|_1$.

Lemma 6.14. $|\hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}|_1 = O(\frac{C_F \hat{m} \hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')}{C_A \hat{T}})$

Proof. One can see that we can bound $|\hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}|_1$ by bounding the total (additive) change of the flow \mathbf{f}^t on all auxiliary arcs during the whole execution of $\hat{\theta}$ -improvement procedure. Furthermore, as the flow \mathbf{f}^t changes only during progress steps, and there is at most $\hat{\theta}^{-2} = \hat{m}^{2\eta}$ of them, it suffices to prove that in each progress step this change is at most $O(C_F \hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') \frac{\hat{m}^{1-2\eta}}{C_A \hat{T}})$.

Now, by Theorem 5.5 and Lemma 6.12, this (additive) change at step t can be bounded as

$$\sum_{e \in S} |\bar{\kappa}_e^t| f_e^t \leq C_F F_A \sum_{e \in S} |\bar{\kappa}_e^t| \leq 4C_F F_A \left(\sum_{e \in S} \delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e + \sum_{e \in S} \hat{\kappa}_e^t \right), \quad (29)$$

where S is the set of auxiliary arcs. By Cauchy-Schwarz inequality, we get that

$$\sum_{e \in S} \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e \leq \sum_{e \in S} \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e \leq \sqrt{\|\boldsymbol{\rho}(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\boldsymbol{\nu}^t, 2}^2 \boldsymbol{\nu}^t(S)} \leq O(\hat{m}),$$

where we used (16) and the fact that by Lemma 6.12 $\boldsymbol{\nu}^t(S)$ is $O(\hat{m})$. Similarly, we obtain that

$$\sum_{e \in S} \hat{\kappa}_e^t \leq \sum_{e \in S} \nu_e^t \hat{\kappa}_e^t \leq \sqrt{\|\hat{\boldsymbol{\kappa}}^t\|_{\boldsymbol{\nu}^t, 2}^2 \boldsymbol{\nu}^t(S)} \leq O(\sqrt{\hat{m}}),$$

where we used that fact that $\|\hat{\kappa}^t\|_{\nu^t, 2} \leq \frac{1}{16}$.

Plugging the above to bounds back into (29) and recalling that we always set $\delta^t := (2C_\delta \hat{\theta} \sqrt{\widehat{m}})^{-1}$, we obtain that

$$\begin{aligned} |\hat{\boldsymbol{\sigma}} - \bar{\boldsymbol{\sigma}}|_1 &\leq 4C_{FFA} \left(\sum_{e \in S} \delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e + \sum_{e \in S} \hat{\kappa}_e^t \right) \leq O\left(\frac{C_{FFA} \sqrt{\widehat{m}}}{C_\delta \hat{\theta}}\right) \\ &\leq O\left(\frac{C_F \hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0}) \widehat{m}^{1-2\eta}}{C_A \widehat{T}}\right) \leq O\left(\frac{C_F \hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') \widehat{m}^{1-2\eta}}{C_A \widehat{T}}\right), \end{aligned}$$

where we utilized (28), as well as, the fact that, due to our stopping condition for $\hat{\theta}$ -improvement phase, we can always assume that $\hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0})$ is $O(\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}'))$. The lemma follows. \square

Clearly, by setting C_A to be a sufficiently large constant, we can ensure that the total measure increase due to fixing procedure will not lead to violation of Invariant 5.2.

$\hat{\theta}$ -smoothness on Light Arcs

As the final step of our analysis, we prove now that in the course of our algorithm – after the modifications described above – all the electrical flows $\hat{\mathbf{f}}^t$ that we compute are indeed $\hat{\theta}$ -smooth on the set of light arcs. That is, the assumption underlying the analysis performed in Section 6.1 is indeed justified.

To this end, let us fix some $\hat{\boldsymbol{\sigma}}$ -feasible and $\hat{\gamma}$ -centered solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ in our augmented graph \bar{G} and let $\hat{\mathbf{f}}^t$ be the associated electrical $\hat{\boldsymbol{\sigma}}$ -flow. For convenience, we drop from now on all the references to t in our notation.

Our proof will take advantage of the dual nature of electrical flows. In particular, it will be instrumental for us to consider the vertex potentials ϕ that induce the electrical flow $\hat{\mathbf{f}}$ via (5). The crucial property of these potentials is that they provide an embedding of all the vertices of \bar{G} into a line. To make it precise, for a given arc $e = (u, v)$, let us denote by ϕ_e^- (resp. ϕ_e^+): the value of ϕ_u (resp. ϕ_v), if $\phi_u \leq \phi_v$; and the value of ϕ_v (resp. ϕ_u), otherwise. In other words, ϕ_e^- (resp. ϕ_e^+) is the coordinate of the left-most (resp. right-most) endpoint of e in this line embedding.

Observe that by (5) and definition of resistances \mathbf{r} (cf. (14)), we have that for a given arc $e = (u, v)$, the distance Δ_e between the embeddings of its endpoints is

$$\Delta_e := \phi_e^+ - \phi_e^- = |\phi_u - \phi_v| = |\hat{f}_e| r_e = \frac{\mu_e |\hat{f}_e|}{f_e^2} = \frac{\mu_e \rho(\hat{\mathbf{f}}, \mathbf{f})_e}{f_e},$$

and thus by Fact 5.3

$$(1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \frac{\rho(\hat{\mathbf{f}}, \mathbf{f})_e}{f_e} \leq \frac{\Delta_e}{\nu_e} \leq (1 + \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \frac{\rho(\hat{\mathbf{f}}, \mathbf{f})_e}{f_e}. \quad (30)$$

Furthermore, for two subsets $T, U \subseteq \bar{V}$ of vertices of \bar{G} , let us define the distance $\text{dist}(T, U)$ between these sets to be

$$\text{dist}(T, U) := \min_{v \in T, u \in U} |\phi_v - \phi_u|. \quad (31)$$

Also, let us call two such subsets $T \subseteq \bar{V}$ and $U \subseteq \bar{V}$, (Δ, k) -separated, for some $\Delta > 0$ and integer $k \geq 0$, if $\text{dist}(T, U) \geq \Delta$ and $\min\{\mathbf{a}(T), \mathbf{a}(U)\} \geq k$, where $\mathbf{a}(U') := \sum_{v \in U'} a_v$ and a_v is defined in (27).

Now, assume for the sake of contradiction that $\hat{\mathbf{f}}$ is not $\hat{\theta}$ -smooth on the set of light arcs, i.e., there exists an $l^* \leq \log \hat{\theta}^{-3}$ such that

$$\nu(S_{l^*}(\hat{\mathbf{f}}) \setminus E_H^t) = \nu(S^*) > \lfloor \hat{\theta}^3 2^{3l^*} \rfloor, \quad (32)$$

where S^* denotes $S_{l^*}(\hat{\mathbf{f}}) \setminus E_H^t$ with the set $S_{l^*}(\hat{\mathbf{f}})$ defined by (17) and E_H^t denotes the set of heavy arcs. Our main goal is to show that in this case there exist two subsets $T, U \subseteq \bar{V}$ of vertices that are (Δ^*, k^*) -separated with

$$\Delta^* := \frac{C_H \hat{m}^{3\eta}}{14 \cdot 2^{l^*}} \quad \text{and} \quad k^* := \frac{2^{2l^*} \hat{m}^{1-6\eta} (\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}))^2}{C_K F_A^2}, \quad (33)$$

where C_H is the constant from Definition 6.7, F_A is given by (28), and C_K is a sufficiently large constant that does not depend on C_H and will be set later.

To motivate this goal, we prove the following lemma.

Lemma 6.15. *If there exist $T, U \subseteq \bar{V}$ that are (Δ^*, k^*) -separated then*

$$\mathcal{E}_r(\hat{\mathbf{f}}) > C_E \hat{m} \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}),$$

provided C_H is chosen to be large enough.

Observe that the conclusion of this lemma violates the bound from Lemma 6.8. Thus, the resulting contradiction would allow us to conclude that $\hat{\mathbf{f}}$ indeed needs to be $\hat{\theta}$ -smooth on the set of light arcs, as we wanted to prove.

Proof. Note that as $\text{dist}(T, U) \geq \Delta^*$, it must be the case that either $\text{dist}(\{\bar{v}\}, U) \geq \frac{\Delta^*}{2}$ or $\text{dist}(\{\bar{v}\}, T) \geq \frac{\Delta^*}{2}$. (Recall that \bar{v} is the special vertex of \bar{G} that is adjacent to all the auxiliary arcs.) Let us assume – without loss of generality – that the first case holds.

Now, as $\min\{\mathbf{a}(T), \mathbf{a}(U)\} \geq k^*$, we know that, in particular, $\mathbf{a}(U) \geq k^*$. This, in turn, means that at least k^* of auxiliary arcs e must have $\Delta_e \geq \frac{\Delta^*}{2}$. Furthermore, by Lemma 6.12, we know that all but $O(\hat{m}^{8\eta})$ of these arcs have measure 1. So, as k^* is $\tilde{\Omega}(\hat{m}^{1-6\eta})$, by ensuring that the constant C_η in the definition of η ((19)) is big enough, we can conclude that the set \hat{S} of auxiliary arcs with $\Delta_e \geq \frac{\Delta^*}{2}$ and $\nu_e = 1$ has size of at least $\frac{k^*}{2}$.

So, by (30) and Lemma 6.12, we have that, for any such arc e in \hat{S} ,

$$\rho(\hat{\mathbf{f}}, \mathbf{f})_e \geq \frac{\Delta_e f_e}{(1 + \hat{\gamma}) \nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \geq \frac{\Delta_e F_A}{(1 + \hat{\gamma}) C_F \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \geq \frac{\Delta^* F_A}{4 C_F \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})},$$

where we used the fact that $\Delta_e \geq \frac{\Delta^*}{2}$ and $\nu_e = 1$, for all e in \hat{S} .

Now, the above inequality enables us to lowerbound the energy $\mathcal{E}_r(\hat{\mathbf{f}})$ of the flow $\hat{\mathbf{f}}$ using sole contribution of arcs in \hat{S} . We get that

$$\begin{aligned} \mathcal{E}_r(\hat{\mathbf{f}}) &\geq \sum_{e \in \hat{S}} r_e \hat{f}_e^2 \geq \sum_{e \in \hat{S}} (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \rho(\hat{\mathbf{f}}, \mathbf{f})_e^2 \geq (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) |\hat{S}| \left(\frac{\Delta^* F_A}{4 C_F \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \right)^2 \\ &\geq \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) k^* \left(\frac{\Delta^* F_A}{8 C_F \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \right)^2 \geq \Omega \left(\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \hat{m}^{1-6\eta} 2^{2l^*} \left(\frac{C_H^2 \hat{m}^{6\eta}}{C_F^2 C_K 2^{2l^*}} \right) \right) \\ &\geq \Omega \left(\hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \hat{m} \left(\frac{C_H^2}{C_F^2 C_K} \right) \right), \end{aligned}$$

where we used the definition of \mathbf{r} (14) and Fact 5.3.

So, once C_H is chosen to be large enough constant – which we can always ensure to be the case – the lemma follows. (Note that at this point the constant C_F is fixed already and we will make sure that when we later set the constant C_K , it does not depend on the value of C_H .) \square

Finding the (Δ^*, k^*) -separated Sets

In the light of the above, it remains to establish how condition (32) implies the existence of such (Δ^*, k^*) -separated sets T^* and U^* . To this end, for a given $x \in \mathbb{R}$, let us define V_x^- (resp. V_x^+) to be the set of vertices v with $\phi_v \leq x$ (resp. $\phi_v \geq x$). Also, let E_x denote the set of arcs e of \bar{G} such that $\phi_e^- \leq x \leq \phi_e^+$.

Now, let x^* be the smallest x such that $\mathbf{a}(V_x^-) \geq k^*$. If $\mathbf{a}(V_{x^*+\Delta^*}) \geq k^*$ then taking $T^* := V_{x^*}^-$ and $U^* = V_{x^*+\Delta^*}^+$ will clearly constitute the (Δ^*, k^*) -separated sets we are looking for.

So, we can focus on the case that $\mathbf{a}(V_{x^*+\Delta^*}) < k^*$. Let us then take $T^* := V_{x^*}^+ \cap V_{x^*+\Delta^*}^-$. Note that, as $3k^*$ is smaller than the number of all auxiliary arcs, we need to have $\mathbf{a}(T^*) \geq k^*$. Next, let us take $U^* := V_{x^*-\Delta^*}^- \cup V_{x^*+2\Delta^*}^+$. Clearly, $\text{dist}(T^*, U^*) \geq \Delta^*$. Therefore, once we show that $\mathbf{a}(U^*) \geq k^*$, T^* and U^* will constitute the desired (Δ^*, k^*) -separated sets.

We proceed now to showing that indeed $\mathbf{a}(U^*) \geq k^*$. Let us define $F(x) := \sum_{e \in E_x} f_e$ (resp. $\hat{F}(x) := \sum_{e \in E_x} |\hat{f}_e|$) to be the total flow of \mathbf{f} (resp. $\hat{\mathbf{f}}$) flowing through the arcs in E_x . We will be interested in two quantities

$$A^* := \int_{\mathbb{R} \setminus I^*} F(x) dx \quad \text{and} \quad \hat{A} := \int_{\mathbb{R} \setminus I^*} \hat{F}(x) dx,$$

where I^* is an interval $[x^* - \Delta^*, x^* + 2\Delta^*]$. (Observe that if the interval I^* was not excluded, \hat{A} would be equal to the energy $\mathcal{E}_r(\hat{\mathbf{f}})$ of the flow $\hat{\mathbf{f}}$.)

Lowerbounding \hat{A} First, we want to lowerbound \hat{A} . To this end, we note that by (30), for any $e \in S^*$ (recall that $S^* := S_{l^*}(\hat{\mathbf{f}}) \setminus E_H^t$ and thus, in particular, is contain only light arcs), we have that

$$\begin{aligned} \Delta_e &\geq (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \nu_e \frac{\rho(\hat{\mathbf{f}}, \mathbf{f})_e}{f_e} \geq (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \frac{\rho(\hat{\mathbf{f}}, \mathbf{f})_e}{F_H} \\ &\geq (1 - \hat{\gamma}) C_H \frac{\rho(\hat{\mathbf{f}}, \mathbf{f})_e}{\hat{m}^{\frac{1}{2}-3\eta}} \geq (1 - \hat{\gamma}) C_H \frac{\hat{m}^{3\eta}}{2^{l^*+1}} \geq 6\Delta^*, \end{aligned}$$

where we also used (17), (33), and Definition 6.7.

As the interval I^* has length $3\Delta^*$, this means that for any arc $e \in S^*$, the interval $[\phi_e^-, \phi_e^+] \setminus I^*$ has length of at least

$$\Delta_e - 3\Delta^* \geq \frac{\Delta_e}{2}.$$

This, in turn, implies that even if we account for contributions of the arcs from S^* only, we have

that

$$\begin{aligned}
\widehat{A} &= \int_{\mathbb{R} \setminus I^*} \widehat{F}(x) dx \geq \frac{1}{2} \sum_{e \in S^*} \Delta_e |\widehat{f}_e| \geq \frac{(1 - \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{2} \sum_{e \in S^*} \frac{\rho(\widehat{\mathbf{f}}, \mathbf{f})_e \nu_e |\widehat{f}_e|}{f_e} \\
&\geq \frac{(1 - \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{2} \sum_{e \in S^*} \rho(\widehat{\mathbf{f}}, \mathbf{f})_e^2 \nu_e \geq \frac{\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{5} \sum_{e \in S^*} \frac{\nu_e \widehat{m}}{2^{2l^*}} \\
&= \frac{\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{5 \cdot 2^{2l^*}} \boldsymbol{\nu}(S^*) \widehat{m} \geq \frac{\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{5} \widehat{\theta}^3 2^{l^*} \widehat{m} = \frac{\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{5} 2^{l^*} \widehat{m}^{1-3\eta},
\end{aligned} \tag{34}$$

where we used (17), (30), and (32).

Upperbounding A^* Now, we want to upperbound the value of A^* . To do that, let us define \bar{S} to be the set of arcs that have at least one endpoint outside of the interval I^* . Note that by our way of setting up the auxiliary arcs and the fact that by Lemma 6.12 and (25), the total increase of measure of arcs during the $\widehat{\theta}$ -improvement phase is $\widetilde{O}(\widehat{m}^{8\eta})$, we have that

$$\mathbf{a}(U^*) \geq \frac{\boldsymbol{\nu}(\bar{S})}{3} - \widetilde{O}(\widehat{m}^{8\eta}). \tag{35}$$

So, if we are able to show that $\boldsymbol{\nu}(\bar{S}) \geq 4k^* = \widetilde{\Omega}(\widehat{m}^{1-6\eta})$ and ensure again that the constant C_η in definition of η (19) is large enough, we will prove that $\mathbf{a}(U^*) \geq k^*$, as desired.

To establish such lowerbound on $\boldsymbol{\nu}(\bar{S})$, we use (30) and observe that

$$A^* = \int_{\mathbb{R} \setminus I^*} F(x) dx \leq \sum_{e \in \bar{S}} \Delta_e f_e \leq (1 + \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \sum_{e \in \bar{S}} \nu_e \rho(\widehat{\mathbf{f}}, \mathbf{f})_e,$$

where we noted that the only arcs that can contribute to A^* are all in the set \bar{S} . Therefore, by Fact 5.3 and Cauchy-Schwarz inequality, we have that

$$(1 + \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \sum_{e \in \bar{S}} \nu_e \rho(\widehat{\mathbf{f}}, \mathbf{f})_e \leq (1 + \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \sqrt{\left(\sum_{e \in \bar{S}} \nu_e \rho(\widehat{\mathbf{f}}, \mathbf{f})_e^2 \right) \boldsymbol{\nu}(\bar{S})}.$$

So, putting the above two bounds together, we get that

$$A^* \leq (1 + \widehat{\gamma}) \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \sum_{e \in \bar{S}} \nu_e \rho(\widehat{\mathbf{f}}, \mathbf{f})_e \leq 5 \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \sqrt{\widehat{m} \boldsymbol{\nu}(\bar{S})}, \tag{36}$$

where we also used (16) and Lemma 5.7.

At this point, our last needed observation is captured by the following lemma.

Lemma 6.16. *For any $x \in \mathbb{R}$, we have that $F(x) \geq \widehat{F}(x)$.*

Notice that once the above lemma is established, we have that

$$A^* = \int_{\mathbb{R} \setminus I^*} F(x) dx \geq \int_{\mathbb{R} \setminus I^*} \widehat{F}(x) dx = \widehat{A},$$

and, as a result, we can put (34) and (36) together to obtain

$$\begin{aligned} \nu(\bar{S}) &\geq \frac{1}{\widehat{m}} \left(\frac{A^*}{5\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \right)^2 \geq \frac{1}{\widehat{m}} \left(\frac{\widehat{A}}{5\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})} \right)^2 \geq \frac{1}{\widehat{m}} \left(\frac{2^{l^*} \widehat{m}^{1-3\eta}}{25} \right)^2 \\ &\geq \Omega(2^{2l^*} \widehat{m}^{1-6\eta}) \geq 4k^*, \end{aligned}$$

once the constant C_K in the definition (33) of k^* is taken to be large enough. (Note that the term $\frac{\widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{F_A}$ in the definition of k^* (33) is bounded by a constant that is independent of C_H . So, indeed C_K does not depend on C_H , as we wanted to ensure.)

Therefore, by (35), the above bounds shows that indeed $\mathbf{a}(U^*) \geq k^*$, as needed.

At this point, we just need to perform the remaining proof of the lemma and the analysis of our improved algorithm will be concluded.

Proof. The simple, but fundamental, observation we need to make here is that the flow $\widehat{\mathbf{f}}$ – being an electrical $\widehat{\boldsymbol{\sigma}}$ -flow induced by vertex potentials ϕ via relationship (5) – is always flowing in one direction, i.e., from left to right, with respect to the line embedding given by ϕ . This, together with the fact that $\widehat{\mathbf{f}}$ is a $\widehat{\boldsymbol{\sigma}}$ -flow, implies that

$$\sum_{v \in V_x^+} \widehat{\sigma}_v = \widehat{F}(x).$$

On the other hand, \mathbf{f} is also a feasible $\widehat{\boldsymbol{\sigma}}$ -flow, which means that the net inflow into V_x^+ of \mathbf{f} has to be at least $\sum_{v \in V_x^+} \widehat{\sigma}_v$. This gives us that

$$F(x) \geq \sum_{v \in V_x^+} \widehat{\sigma}_v = \widehat{F}(x),$$

as we wanted to establish. □

7 Electrical Flows and the Central Path

In this section, we describe how we can use electrical flows to advance our solution along the central path. In other words, we describe and analyze the implementation of the improvement step and thus prove Theorem 5.5. This implementation is directly inspired by – and, in fact, can be seen as a reinterpretation of – the improvement steps used in path-following method.

Recall that in the improvement step, we are given a $\widehat{\gamma}$ -centered $\widehat{\boldsymbol{\sigma}}$ -feasible solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and our goal is to compute, in $\widetilde{O}(\widehat{m})$ time, a $\widehat{\gamma}$ -centered $\widehat{\boldsymbol{\sigma}}$ -feasible solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ with

$$\widehat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \leq (1 - \delta^t) \widehat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t). \quad (37)$$

We perform this improvement in two main steps. The first one – the descent step – uses the electrical flow $\widehat{\mathbf{f}}^t$ associated with $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and the corresponding vertex potentials $\widehat{\phi}^t$ that induce it, to perform a primal and dual update that results in a new, intermediate, solution $(\widehat{\mathbf{f}}^t, \widehat{\mathbf{s}}^t, \widehat{\boldsymbol{\nu}}^t)$. This intermediate solution is $\widehat{\boldsymbol{\sigma}}$ -feasible and has $\widehat{\mu}(\widehat{\mathbf{f}}^t, \widehat{\mathbf{s}}^t, \widehat{\boldsymbol{\nu}}^t) \leq (1 - \delta^t) \widehat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ as desired, but it might be not $\widehat{\gamma}$ -centered anymore. To fix that, in the second – centering – step, we compute the desired solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ out of $(\widehat{\mathbf{f}}^t, \widehat{\mathbf{s}}^t, \widehat{\boldsymbol{\nu}}^t)$ by using another electrical flow computation that again provides a primal and dual update.

We describe and analyze both of these steps below. Note that as each of these two steps requires only one computation of electrical flow, it can be easily implemented to run in $\widetilde{O}(\widehat{m})$ time, as needed.

Descent Step Let $\hat{\mathbf{f}}^t$ be the electrical $\hat{\sigma}$ -flow associated with the solution $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ and let $\hat{\phi}^t$ be the vertex potentials that induce $\hat{\mathbf{f}}^t$. Consider a new primal-dual solution $(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$ given by

$$\bar{f}_e^t := (1 - \delta^t)f_e^t + \delta^t \hat{f}_e^t \quad (38)$$

$$\bar{s}_e^t := s_e^t - \frac{\delta^t}{(1 - \delta^t)}(\hat{\phi}_u^t - \hat{\phi}_v^t) = s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)f_e^t} \hat{f}_e^t \quad (39)$$

$$\bar{\nu}_e^t := \nu_e^t, \quad (40)$$

for each arc $e = (v, u)$ in \hat{G} , where δ^t satisfies conditions of the theorem and we also used the definition (14) of the resistances that determine $\hat{\mathbf{f}}^t$, as well as, the relationship (5) between electrical flow and the vertex potentials that induce it.

Observe that as $\bar{\mathbf{f}}^t$ is a convex combination of two $\hat{\sigma}$ -flows – the flows \mathbf{f}^t and $\hat{\mathbf{f}}^t$ – it also is an $\hat{\sigma}$ -flow. Furthermore, as all $\nu_e^t \geq 1$, we have $\|\rho(\bar{\mathbf{f}}^t, \mathbf{f}^t)\|_{\nu^t, 4} \geq \|\rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_\infty$ and thus, for each arc e ,

$$\bar{f}_e^t = (1 - \delta^t)f_e^t + \delta^t \hat{f}_e^t \geq (1 - \delta^t)f_e^t - \delta^t |\hat{f}_e^t| = (1 - \delta^t - \delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e) f_e^t \geq (1 - \frac{1}{2} - \sqrt{\hat{\gamma}}) f_e^t > 0, \quad (41)$$

and similarly

$$\bar{s}_e^t = s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)f_e^t} \hat{f}_e^t \geq s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)f_e^t} |\hat{f}_e^t| = s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)} \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e \geq (1 - 2\sqrt{\hat{\gamma}}) s_e^t > 0.$$

So, $(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$ is $\hat{\sigma}$ -feasible, as desired.

Let us now analyze the value of $\hat{\mu}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$. To this end, observe that, for any arc e ,

$$\begin{aligned} \hat{\mu}_e^t &= \frac{\bar{f}_e^t \bar{s}_e^t}{\bar{\nu}_e^t} = (\nu_e^t)^{-1} ((1 - \delta^t)f_e^t + \delta^t \hat{f}_e^t) (s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)f_e^t} \hat{f}_e^t) \\ &= (\nu_e^t)^{-1} \left((1 - \delta^t)f_e^t s_e^t + \delta^t \hat{f}_e^t s_e^t - \delta^t \frac{s_e^t}{(1 - \delta^t)f_e^t} \hat{f}_e^t (1 - \delta^t)f_e^t - (\delta^t)^2 \frac{s_e^t}{(1 - \delta^t)f_e^t} (\hat{f}_e^t)^2 \right) \\ &= (\nu_e^t)^{-1} \left((1 - \delta^t)\mu_e^t + \delta^t \hat{f}_e^t s_e^t - \delta^t \hat{f}_e^t s_e^t - (\delta^t)^2 \frac{s_e^t}{(1 - \delta^t)} f_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^2 \right) \\ &= \left(1 - \delta^t - \frac{(\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e)^2}{(1 - \delta^t)} \right) \hat{\mu}_e^t. \end{aligned} \quad (42)$$

So, by definition (10) and the fact that $\frac{(\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e)^2}{(1 - \delta^t)} \geq 0$ for all e , we see that

$$\hat{\mu}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t) \leq (1 - \delta^t) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$$

and this inequality would be an equality if the second-order terms (i.e., terms quadratic in δ^t) were ignored. (Also, if these terms were not present, the centrality of the solution would be preserved too.)

Finally, let us focus on analyzing the centrality of $(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$. To this end, note that by definition (11) and by (42) above we have

$$\begin{aligned}
\|\hat{\boldsymbol{\mu}}^t - \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)\mathbf{1}\|_{\bar{\boldsymbol{\nu}}^t, 2}^2 &\leq \|\hat{\boldsymbol{\mu}}^t - (1 - \delta^t)\hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)\mathbf{1}\|_{\bar{\boldsymbol{\nu}}^t, 2}^2 \leq \sum_e \bar{\nu}_e^t (\hat{\mu}_e^t - (1 - \delta^t)\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t))^2 \\
&= \sum_e \nu_e^t \left((1 - \delta^t)(\hat{\mu}_e^t - \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)) - \frac{(\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e)^2}{(1 - \delta^t)} \hat{\mu}_e^t \right)^2 \\
&\leq 2 \left((1 - \delta^t)^2 \sum_e \nu_e^t (\hat{\mu}_e^t - \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t))^2 + \sum_e \nu_e^t \frac{(\delta^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e)^4}{(1 - \delta^t)^2} (\hat{\mu}_e^t)^2 \right) \quad (43) \\
&\leq 2 \left((1 - \delta^t)^2 \|\hat{\boldsymbol{\mu}}^t - \hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)\mathbf{1}\|_{\boldsymbol{\nu}^t, 2}^2 + \frac{(1 + \hat{\gamma})^2 \hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)^2}{(1 - \delta^t)^2} (\delta^t)^4 \sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^4 \right) \\
&\leq 2 \left((1 - \delta^t)^2 \hat{\gamma}^2 \hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)^2 + \frac{(1 + \hat{\gamma})^2 \hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)^2}{(1 - \delta^t)^2} (\delta^t)^4 \sum_e \nu_e^t \rho(\hat{\mathbf{f}}^t, \mathbf{f}^t)_e^4 \right) \\
&\leq 2 \left((1 - \delta^t)^2 \hat{\gamma}^2 + \frac{(1 + \hat{\gamma})^2}{(1 - \delta^t)^2} (\delta^t)^4 \|\boldsymbol{\rho}(\hat{\mathbf{f}}^t, \mathbf{f}^t)\|_{\boldsymbol{\nu}^t, 4}^4 \right) \hat{\boldsymbol{\mu}}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)^2 \\
&\leq 10 \hat{\gamma}^2 \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)^2.
\end{aligned}$$

In the above derivation, the first inequality follows as the $\|\boldsymbol{\mu} - t\mathbf{1}\|_{\boldsymbol{\nu}, 2}$ is always minimized by taking $t = \hat{\boldsymbol{\mu}}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$. We also used the fact that $(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ is $\hat{\gamma}$ -centered, Fact 5.3 and the upperbound on δ^t .

Therefore, we see that the price of making progress on the duality gap is that the centrality of our solution could deteriorate by a factor of at most three.

Centering Step To alleviate this possible increase of centrality, we apply a second step that restores the centrality back within the desired bounds while not increasing the duality gap (so to not to counter the progress on the duality gap that we just achieved).

To this end, consider a flow $\bar{\mathbf{f}}^*$ in \widehat{G} defined as

$$\bar{f}_e^* := \frac{\hat{\mu}_e^t - \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)}{\hat{\mu}_e^t} \bar{f}_e^t, \quad (44)$$

for every arc e of \widehat{G} . Note that the flow $\bar{\mathbf{f}}^*$ might (and actually will) not be feasible in \widehat{G} , as some of \bar{f}_e^* can be negative.

Now, consider a flow $\bar{\mathbf{f}}'$ given by

$$\bar{f}'_e := \bar{f}_e^t - \bar{f}_e^* = \left(1 - \frac{\hat{\mu}_e^t - \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)}{\hat{\mu}_e^t}\right) \bar{f}_e^t = \frac{\hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)}{\hat{\mu}_e^t} \bar{f}_e^t, \quad (45)$$

for each arc e . Observe that $\bar{\mathbf{f}}'$ is feasible in \widehat{G} (i.e., $\bar{f}'_e \geq 0$, for all e) and

$$\frac{\bar{f}'_e \bar{s}_e^t}{\bar{\nu}_e^t} = \frac{\hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)}{\hat{\mu}_e^t \bar{\nu}_e^t} \bar{f}_e^t \bar{s}_e^t = \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t), \quad (46)$$

for each arc e . That is, $(\bar{\mathbf{f}}', \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$ is 0-centered with $\hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}', \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t) = \hat{\boldsymbol{\mu}}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t)$.

So, this solution would be a perfect candidate for $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ except that the flow $\bar{\mathbf{f}}'$ does not need to be a $\hat{\boldsymbol{\sigma}}$ -flow and thus this solution might not be $\hat{\boldsymbol{\sigma}}$ -feasible.

To fix that – and obtain our desired solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ – let $\tilde{\boldsymbol{\sigma}}$ be the demand vector of the flow $\bar{\mathbf{f}}^*$, and consider an electrical $\tilde{\boldsymbol{\sigma}}$ -flow $\tilde{\mathbf{f}}^t$ that corresponds to resistances

$$\tilde{r}_e^t := \frac{\bar{s}_e^t}{\bar{f}'_e}, \quad (47)$$

for each arc e and let $\tilde{\phi}$ be the corresponding vertex potentials.

Let us define $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ to be

$$\begin{aligned} f_e^{t+1} &:= \bar{f}'_e + \tilde{f}_e^t \\ s_e^{t+1} &:= \bar{s}_e^t - (\tilde{\phi}_u^t - \tilde{\phi}_v^t) = \bar{s}_e^t - \frac{\bar{s}_e^t}{\bar{f}'_e} \tilde{f}_e^t \\ \nu_e^{t+1} &:= \bar{\nu}_e^t, \end{aligned} \quad (48)$$

for each arc $e = (v, u)$.

Clearly, now \mathbf{f}^{t+1} is a $\hat{\boldsymbol{\sigma}}$ -flow, as desired. Let us analyze its centrality. To this end, let us fix some arc e , and notice that

$$\begin{aligned} \hat{\mu}_e^{t+1} &= \frac{f_e^{t+1} s_e^{t+1}}{\nu_e^{t+1}} = \frac{(\bar{f}'_e + \tilde{f}_e^t)(\bar{s}_e^t - \frac{\bar{s}_e^t}{\bar{f}'_e} \tilde{f}_e^t)}{\bar{\nu}_e^t} \\ &= (\bar{\nu}_e^t)^{-1} \left(\bar{f}'_e \bar{s}_e^t + \bar{s}_e^t \tilde{f}_e^t - \bar{s}_e^t \tilde{f}_e^t - \frac{\bar{s}_e^t}{\bar{f}'_e} (\tilde{f}_e^t)^2 \right) \\ &= \hat{\mu}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t) - \frac{\bar{s}_e^t}{\bar{\nu}_e^t \bar{f}'_e} (\tilde{f}_e^t)^2, \end{aligned}$$

where we used (46). So, we see in particular that

$$\hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \leq \hat{\mu}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t) \leq (1 - \delta^t) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t),$$

as needed.

Now, by our derivation above, we have that

$$\|\hat{\boldsymbol{\mu}}^{t+1} - \hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \mathbf{1}\|_{\boldsymbol{\nu}^{t+1}, 2}^2 \leq \|\hat{\boldsymbol{\mu}}^{t+1} - \hat{\mu}(\bar{\mathbf{f}}^t, \bar{\mathbf{s}}^t, \bar{\boldsymbol{\nu}}^t) \mathbf{1}\|_{\bar{\boldsymbol{\nu}}^t, 2}^2 = \sum_e \bar{\nu}_e^t \left(\frac{\bar{s}_e^t}{\bar{\nu}_e^t \bar{f}'_e} (\tilde{f}_e^t)^2 \right)^2.$$

To bound the resulting expression, let us note that by Cauchy-Schwarz inequality and the fact that measures are always at least 1 we have

$$\sum_e \bar{\nu}_e^t \left(\frac{\bar{s}_e^t}{\bar{\nu}_e^t \bar{f}'_e} (\tilde{f}_e^t)^2 \right)^2 \leq \left(\max_e \frac{\bar{s}_e^t}{\bar{\nu}_e^t \bar{f}'_e} (\tilde{f}_e^t)^2 \right) \left(\sum_e \frac{\bar{s}_e^t}{\bar{f}'_e} (\tilde{f}_e^t)^2 \right) \leq \left(\sum_e \frac{\bar{s}_e^t}{\bar{f}'_e} (\tilde{f}_e^t)^2 \right)^2.$$

Now, the key insight here is that by (47),

$$\left(\sum_e \frac{\bar{s}_e^t}{\bar{f}'_e} (\tilde{f}_e^t)^2 \right)^2 = \left(\mathcal{E}_{\tilde{r}^t}(\bar{\mathbf{f}}^t) \right)^2.$$

So, by bounding the energy of the electrical flow \tilde{f}_e^t we will be able to bound the centrality of our solution $\hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$. To bound this energy, we will first bound the energy $\mathcal{E}_{\tilde{r}^t}(\tilde{\mathbf{f}}^*)$ of the flow $\tilde{\mathbf{f}}^*$ and use the fact that both $\tilde{\mathbf{f}}^*$ and $\tilde{\mathbf{f}}^t$ are $\tilde{\boldsymbol{\sigma}}$ -flows and thus, by definition, $\tilde{\mathbf{f}}^t$ is minimizing the energy among all the $\tilde{\boldsymbol{\sigma}}$ -flows.

Observe that by definition (44) of the flow $\tilde{\mathbf{f}}^*$, the fact that $(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)$ is $3\hat{\gamma}$ -centered – cf. (43) – and Fact 5.3, we have that

$$\begin{aligned}
\mathcal{E}_{\tilde{r}^t}(\tilde{\mathbf{f}}^*) &= \sum_e \frac{\tilde{s}_e^t}{\tilde{f}_e^t} \left(\frac{\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)}{\hat{\mu}_e^t} \tilde{f}_e^t \right)^2 \leq \sum_e \frac{\tilde{\nu}_e^t \hat{\mu}_e^t}{\tilde{f}_e^t} \frac{(\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t))^2 \tilde{f}_e^t}{(1 - 3\hat{\gamma}) \hat{\mu}_e^t \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)} \\
&\leq \frac{1}{(1 - 3\hat{\gamma})} \sum_e \tilde{\nu}_e^t \frac{(\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t))^2 \tilde{f}_e^t}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \tilde{f}_e^t} = \frac{1}{(1 - 3\hat{\gamma})} \sum_e \tilde{\nu}_e^t \frac{(\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t))^2 \hat{\mu}_e^t}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)^2} \\
&\leq \frac{(1 + 3\hat{\gamma})}{(1 - 3\hat{\gamma})} \sum_e \tilde{\nu}_e^t \frac{(\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t))^2}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)} = \frac{(1 + 3\hat{\gamma})}{(1 - 3\hat{\gamma})} \frac{\|\hat{\boldsymbol{\mu}}^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \mathbf{1}\|_{\tilde{\boldsymbol{\nu}}^t, 2}^2}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t)} \quad (49) \\
&\leq \frac{(1 + 3\hat{\gamma})}{(1 - 3\hat{\gamma})} 9\hat{\gamma}^2 \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \leq 10\hat{\gamma}^2 \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \leq 20\hat{\gamma}^2 \hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}),
\end{aligned}$$

where we also used the definition (45) of the flow $\tilde{\mathbf{f}}^t$.

In the light of the above discussion, we can conclude that

$$\begin{aligned}
\|\hat{\boldsymbol{\mu}}^{t+1} - \hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \mathbf{1}\|_{\boldsymbol{\nu}^{t+1}, 2}^2 &\leq \sum_e \tilde{\nu}_e^t \left(\frac{\tilde{s}_e^t}{\tilde{\nu}_e^t \tilde{f}_e^t} (\tilde{f}_e^t)^2 \right)^2 \leq \left(\sum_e \frac{\tilde{s}_e^t}{\tilde{f}_e^t} (\tilde{f}_e^t)^2 \right)^2 \\
&\leq (20\hat{\gamma}^2 \hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}))^2 \leq \hat{\gamma}^2 \hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})^2,
\end{aligned}$$

as $\hat{\gamma} \leq \frac{1}{20}$. So, indeed $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ is $\hat{\gamma}$ -centered.

Now, to prove that $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ is also $\hat{\boldsymbol{\sigma}}$ -feasible, we just need to show that for any arc e ,

$$\rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e = \frac{|\tilde{f}_e^t|}{\tilde{f}_e^t} \leq \frac{1}{2}.$$

To this end, note that by (46) and (49) we have

$$\nu_e^t \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e^2 = \frac{\tilde{s}_e^t \tilde{f}_e^t}{(\tilde{f}_e^t)^2} (\tilde{f}_e^t)^2 = \frac{\tilde{s}_e^t}{\tilde{f}_e^t} (\tilde{f}_e^t)^2 \leq \mathcal{E}_{\tilde{r}^t}(\tilde{\mathbf{f}}^t) \leq \mathcal{E}_{\tilde{r}^t}(\tilde{\mathbf{f}}^*) \leq 10\hat{\gamma}^2 \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t) \leq \frac{1}{1600} \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\boldsymbol{\nu}}^t). \quad (50)$$

Thus, indeed, we can conclude that we obtained a $\hat{\gamma}$ -centered $\hat{\boldsymbol{\sigma}}$ -feasible solution $(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1})$ with $\hat{\mu}(\mathbf{f}^{t+1}, \mathbf{s}^{t+1}, \boldsymbol{\nu}^{t+1}) \leq (1 - \delta^t) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$, as desired.

This concludes the proof of the first part of the Theorem 5.5. The proof of the second part appears in Appendix F.

8 Rounding Fractional Bipartite \mathbf{b} -Matchings

In this section, we show how given a fractional \mathbf{b} -matching \mathbf{x} in some bipartite graph $G = (P \cup Q, E)$ with $m = |E|$ edges, one can find in $\tilde{O}(m)$ time an integral \mathbf{b} -matching \mathbf{x}^* in G whose size is at least $\lfloor \|\mathbf{x}\|_1 \rfloor$. In other words, we prove Theorem 3.3.

Rounding Perfect Matchings

Let us first consider the case when \mathbf{x} is just a fractional perfect matching, i.e., $b_v = 1$ for all vertices and the size $|\mathbf{x}|_1$ of \mathbf{x} is $\frac{|b|_1}{2}$, i.e., the fractional degree of each vertex in \mathbf{x} is 1. We claim that in this case we can just use Theorem 2.4 to obtain an integral perfect matching in $\tilde{O}(m)$ time.

To see why this is the case, consider a $|P| \times |Q|$ matrix $\mathbf{M}^{\mathbf{x}}$ in which rows and columns are indexed by vertices from P and Q , respectively, and the entries are given by $M_{p,q}^{\mathbf{x}} := x_{(p,q)}$ if the edge (p,q) exists in G ; and 0, otherwise. Observe that if \mathbf{x} is perfect and all b_v are equal to 1 then we need to have $|P| = |Q|$. Thus, $\mathbf{M}^{\mathbf{x}}$ is a square matrix. Furthermore, $\mathbf{M}^{\mathbf{x}}$ needs to be also doubly-stochastic, as for any row indexed by vertex $p \in P$ (resp. column indexed by vertex $q \in Q$), the sum $\sum_{q' \in Q} M_{p,q'}^{\mathbf{x}}$ (resp. $\sum_{p' \in P} M_{p',q}^{\mathbf{x}}$) of the entries in this row (resp. column) is equal to $\sum_{e \in E(p)} x_e = b_p = 1$ (resp. $\sum_{e \in E(q)} x_e = b_q = 1$). So, invoking Theorem 2.4, we can obtain in $\tilde{O}(m)$ time an integral matching \mathbf{x}^* in the support of $\mathbf{M}^{\mathbf{x}}$ that is also the support of the edge set E of our graph G .

Rounding Non-Perfect Matchings

Now, to recover the desired integral matching in the case when \mathbf{x} is not necessarily perfect (but still all b_v are equal to 1), our first step is to extend \mathbf{x} to a perfect matching $\bar{\mathbf{x}}$ in a certain augmented graph \bar{G} that is created from G by adding some dummy edges and vertices to it.

More precisely, let d_P (resp. d_Q) be the total deficits of vertices in P (resp. in Q), i.e.,

$$d_P := |P| - \sum_{e \in E(p), p \in P} x_e \quad \text{and} \quad d_Q := |Q| - \sum_{e \in E(q), q \in Q} x_e.$$

Note that the size $|\mathbf{x}|_1$ of \mathbf{x} has to be exactly $|P| - d_P = |Q| - d_Q$. We add to the vertex set Q , $\lceil d_P \rceil$ (resp. to the vertex set P , $\lceil d_Q \rceil$) dummy vertices $\bar{q}_1, \dots, \bar{q}_{\lceil d_P \rceil}$ (resp. $\bar{p}_1, \dots, \bar{p}_{\lceil d_Q \rceil}$). Next, we extend the fractional matching \mathbf{x} to $\bar{\mathbf{x}}$ by going over each non-dummy vertex $p \in P$ (resp. $q \in Q$) and fractionally matching it to the dummy vertices $\bar{q}_1, \dots, \bar{q}_{\lceil d_P \rceil}$ (resp. $\bar{p}_1, \dots, \bar{p}_{\lceil d_Q \rceil}$), so to ensure that its fractional degree becomes 1 and the fractional degree of dummy vertices never exceeds one. It is not hard to see that by employing a simple greedy approach we can achieve this goal in $\tilde{O}(m)$ time and, furthermore, ensure that: (1) each non-dummy vertex is matched to at most two dummy vertices in $\bar{\mathbf{x}}$; (2) at the end, there are at most two dummy vertices, say, $\bar{p}_{\lceil d_Q \rceil}$ and $\bar{q}_{\lceil d_P \rceil}$, (one on each side of the bipartition) that are yet not fully matched in $\bar{\mathbf{x}}$. To alleviate the latter problem, we just match these two dummy vertices to each other (one can check that their deficits have to be equal) and take the set of edges \bar{E} of our augmented graph \bar{G} to be the support of the matching $\bar{\mathbf{x}}$. (Note that by property (1), the size of this support will be still $O(m)$.)

Clearly, $\bar{\mathbf{x}}$ is a perfect matching in \bar{G} , so we can use the $\tilde{O}(m)$ -time procedure we described above to get an integral perfect matching $\bar{\mathbf{x}}^*$ in that graph. Once we do that, we take our desired integral matching \mathbf{x}^* in G to be $\bar{\mathbf{x}}^*$ after we removed from it all the edges of $\bar{\mathbf{x}}^*$ that are not in G , i.e., all the edges that are incident to dummy vertices. Obviously, \mathbf{x}^* is a feasible matching in G and it is integral. To see that its size is at least $\lfloor |\mathbf{x}|_1 \rfloor$, note that, as there is at most $\lceil d_P \rceil + \lceil d_Q \rceil$ dummy vertices in \bar{G} , there could be at most that many edges incident to these vertices in $\bar{\mathbf{x}}^*$. But, as $\bar{\mathbf{x}}^*$ is perfect, its size is equal to

$$\frac{|P| + \lceil d_P \rceil + |Q| + \lceil d_Q \rceil}{2} = \frac{|P| - \lceil d_P \rceil + |Q| - \lceil d_Q \rceil}{2} + \lceil d_P \rceil + \lceil d_Q \rceil = \lfloor |\mathbf{x}|_1 \rfloor + \lceil d_P \rceil + \lceil d_Q \rceil,$$

where we used the fact that $|P| - d_P = |Q| - d_Q = |\mathbf{x}|_1$. Thus, indeed after removing at most $\lceil d_P \rceil + \lceil d_Q \rceil$ edges from $\bar{\mathbf{x}}^*$, the resulting integral matching \mathbf{x}^* will have its size $|\mathbf{x}^*|_1$ to be at least $\lfloor |\mathbf{x}|_1 \rfloor$, as desired.

Rounding \mathbf{b} -Matchings

In the light of the above, it remains to show how to deal with the case when in the demand vector \mathbf{b} there are some b_v that are bigger than 1 (and thus some of the entries of \mathbf{x} could be bigger than 1, as well). To this end, let us observe first that if there is an edge $e = (p, q)$ with $x_e \geq 1$, we can just subtract $\lfloor x_e \rfloor$ copies of this edge from our matching right away, while decreasing the demands b_p and b_q of e 's endpoints accordingly, i.e., by $\lfloor x_e \rfloor$. (Note that by feasibility of \mathbf{x} , $b_p, b_q \geq \lfloor x_e \rfloor$.) So, one can see that if $\bar{\mathbf{x}}$ is the fractional matching \mathbf{x} after we made such transformation and $\bar{\mathbf{b}}$ are the corresponding demands, then once we compute an integral $\bar{\mathbf{b}}$ -matching $\bar{\mathbf{x}}^*$ of size at least $\lfloor \|\bar{\mathbf{x}}\|_1 \rfloor$ from $\bar{\mathbf{x}}$, we can just add back these subtracted $\lfloor x_e \rfloor$ copies of edge e to $\bar{\mathbf{x}}^*$ to obtain the desired integral \mathbf{b} -matching \mathbf{x}^* of size at least $\lfloor \|\bar{\mathbf{x}}\|_1 \rfloor + \lfloor x_e \rfloor = \lfloor \|\mathbf{x}\|_1 \rfloor$.

Therefore, we can assume from now on that in our \mathbf{b} -matching \mathbf{x} all x_e s are smaller than one (but still we can have some demands b_v to be bigger than one). To round such fractional \mathbf{b} -matchings, for each vertex $v \in V$ that has its demand b_v bigger than 1, we split it into b_v vertices v^1, \dots, v^{b_v} – each with demand one. Next, for every edge e that was previously incident to v , we connect it to the new vertices and distribute its fractional weight x_e in \mathbf{x} among these new vertices. Again, by applying a simple greedy approach we can ensure that each edge is connected to at most two among the vertices v^1, \dots, v^{b_v} and none of these vertices has its fractional degree bigger than 1. (Note that this means, in particular, that once we apply such splitting to all vertices with $b_v > 1$ then the support of the corresponding “split” fractional matching is at most by a factor of four larger than the support of \mathbf{x} .) Clearly, at this point, we are again in situation where we just need to round a fractional bipartite matching (with all demands being at most 1). Thus, we can use our rounding procedure we described above and recover the integral matching we are seeking. This finishes the proof of Theorem 3.3.

Acknowledgments. We are grateful to Andrew Goldberg, Jonathan Kelner, Lap Chi Lau, Gary Miller, Richard Peng, Seth Pettie, Daniel Spielman, and Shang-Hua Teng for a number of helpful discussions on this topic. We also thank Monika Henzinger, Satish Rao, and Jens Vygen for useful feedback on the manuscript.

References

- [ABMP91] H. Alt, N. Blum, K. Mehlhorn, and M. Paul. Computing a maximum cardinality matching in a bipartite graph in time $O(n^{1.5} \sqrt{m/\log n})$. *Inf. Process. Lett.*, 37(4):237–240, 1991.
- [AHK12] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [AHU74] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Longman Publishing Co., Boston, MA, USA, 1st edition, 1974.

- [AMO93] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, 1993.
- [AMOR95] R. K. Ahuja, T. L. Magnanti, J. B. Orlin, and M. R. Reddy. *Applications of Network Optimization*, volume 7 of *Handbooks in Operations Research and Management Science*. North-Holland, 1995.
- [BH74] J. R. Bunch and J. E. Hopcroft. Triangular factorization and inversion by fast matrix multiplication. *Mathematics of Computation*, 28(125):231–236, 1974.
- [Bol98] B. Bollobas. *Modern Graph Theory*. Springer, 1998.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [CKM⁺11] P. Christiano, J. Kelner, A. Mądry, D. Spielman, and S.-H. Teng. Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs. In *STOC'11: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 273–281, 2011.
- [CLRS09] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.
- [CW90] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9:251–280, 1990.
- [DS08] S. I. Daitch and D. A. Spielman. Faster approximate lossy generalized flow via interior point algorithms. In *STOC'08: Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 451–460, 2008.
- [Edm65] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [EFS56] P. Elias, A. Feinstein, and C. E. Shannon. A note on the maximum flow through a network. *IRE Transactions on Information Theory*, 2, 1956.
- [Ege31] J. Egerváry. Matrixok kombinatorius tulajdonságairól. *Matematikai és Fizikai Lapok*, 38:16–28, 1931.
- [ET75] S. Even and R. E. Tarjan. Network flow and testing graph connectivity. *SIAM Journal on Computing*, 4(4):507–518, 1975.
- [FF56] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.
- [FM95] T. Feder and R. Motwani. Clique partitions, graph compression and speeding-up algorithms. *Journal of Computer and System Sciences*, 51(2):261–272, 1995.
- [Fro12] F. G. Frobenius. Über matrizen aus nicht negativen elementen. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 456–477, 1912.
- [Fro17] F. G. Frobenius. Über zerlegbare determinanten. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 274–277, 1917.

- [GK04] A. V. Goldberg and A. V. Karzanov. Maximum skew-symmetric flows and matchings. *Mathematical Programming*, 100(3):537–568, 2004.
- [GKK10] A. Goel, M. Kapralov, and S. Khanna. Perfect matchings in $O(n \log n)$ time in regular bipartite graphs. In *STOC'10: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing*, pages 39–46, 2010.
- [GR98] A. V. Goldberg and S. Rao. Beyond the flow decomposition barrier. *Journal of the ACM*, 45(5):783–797, 1998.
- [GT91] H. N. Gabow and R. E. Tarjan. Faster scaling algorithms for general graph matching problems. *Journal of the ACM*, 38(4):815–853, 1991.
- [Har09] N. J. A. Harvey. Algebraic algorithms for matching and matroid problems. *SIAM Journal on Computing*, 39(2):679–702, 2009.
- [HK73] J. Hopcroft and R. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.
- [Hof60] A. J. Hoffman. Some recent applications of the theory of linear inequalities to extremal combinatorial analysis. In *Proceedings of Symposia in Applied Mathematics*, volume 10, pages 113–127, 1960.
- [Kar73] A. V. Karzanov. O nakhozhenii maksimal'nogo potoka v setyakh spetsial'nogo vida i nekotorykh prilozheniyakh. *Matematicheskie Voprosy Upravleniya Proizvodstvom*, 5:81–94, 1973. (in Russian; title translation: On finding maximum flows in networks with special structure and some applications).
- [KKL13] S. Khanna, T. C. Kwok, and L. C. Lau, 2013. Personal communication.
- [KLOS13] J. A. Kelner, Y. T. Lee, L. Orecchia, and A. Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. 2013.
- [KLOS14] J. A. Kelner, Y. T. Lee, L. Orecchia, and A. Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *SODA'14: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [KMP10] I. Koutis, G. L. Miller, and R. Peng. Approaching optimality for solving SDD systems. In *FOCS'10: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 235–244, 2010.
- [KMP11] I. Koutis, G. L. Miller, and R. Peng. A nearly $m \log n$ -time solver for SDD linear systems. In *FOCS'11: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 590–598, 2011.
- [Kön15] D. König. Vonalrendszerek és determinánsok. *Mathematikai és Természettudományi Értesítő*, 33:221–229, 1915.

- [Kön16] D. König. Über graphen und ihre anwendung auf determinantentheorie und mengenlehre. *Mathematische Annalen*, 77:453–465, 1916.
- [Kön23] D. König. Sur un problème de la théorie générale des ensembles et la théorie des graphes. *Revue de Métaphysique et de Morale*, 30:443–449, 1923.
- [Kön31] D. König. Graphok és matrixok. *Matematikai és Fizikai Lapok*, 38:116–119, 1931.
- [KOSZ13] J. A. Kelner, L. Orecchia, A. Sidford, and Z. A. Zhu. A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In *STOC'13: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing*, pages 911–920, 2013.
- [KRT94] V. King, S. Rao, and R. Tarjan. A faster deterministic maximum flow algorithm. *Journal of Algorithms*, 17(3):447–474, 1994.
- [Lau13] L. C. Lau, 2013. Personal communication.
- [Lov79] L. Lovász. On determinants, matchings and random algorithms. *Fundamentals of Computation Theory*, 565–574, 1979.
- [LP86] L. Lovász and D. M. Plummer. *Matching Theory*. Elsevier Science, 1986.
- [LRS13] Y. T. Lee, S. Rao, and N. Srivastava. A new approach to computing maximum flows using electrical flows. In *STOC'13: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing*, pages 755–764, 2013.
- [Mađ10] A. Mađry. Fast approximation algorithms for cut-based problems in undirected graphs. In *FOCS'10: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 245–254, 2010.
- [Mađ11] A. Mađry. *From Graphs to Matrices, and Back: New Techniques for Graph Algorithms*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [MS04] M. Mucha and P. Sankowski. Maximum matchings via Gaussian elimination. In *FOCS'04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 248–255, 2004.
- [Muc05] M. Mucha. *Finding maximum matchings via Gaussian elimination*. PhD thesis, University of Warsaw, 2005.
- [MV80] S. Micali and V. V. Vazirani. An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In *FOCS'80: Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science*, pages 17–27, 1980.
- [Orl13] J. B. Orlin. Max flows in $O(nm)$ time, or better. In *STOC'13: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing*, pages 765–774, 2013.
- [RV89] M. O. Rabin and V. V. Vazirani. Maximum matchings in general graphs through randomization. *J. Algorithms*, 10(4):557–567, December 1989.
- [Sch02] A. Schrijver. On the history of the transportation and maximum flow problems. *Mathematical Programming*, 91:437–445, 2002.

- [Sch03] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.
- [Sch05] A. Schrijver. On the history of combinatorial optimization (till 1960). In R. W. K. Aardal, G.L. Nemhauser, editor, *Handbook of Discrete Optimization*, pages 1–68. Elsevier, 2005.
- [She09] J. Sherman. Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$ -approximations to sparsest cuts. In *FOCS'09: Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 363–372, 2009.
- [She13] J. Sherman. Nearly maximum flows in nearly linear time. In *FOCS'13: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, 2013.
- [ST03] D. A. Spielman and S.-H. Teng. Solving sparse, symmetric, diagonally-dominant linear systems in time $O(m^{1.31})$. In *FOCS'03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 416–427, 2003.
- [ST04] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *STOC'04: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing*, pages 81–90, 2004.
- [Tut47] W. T. Tutte. The factorization of linear graphs. *Journal of the London Mathematical Society*, 22:107–111, 1947.
- [Vaz94] V. V. Vazirani. A theory of alternating paths and blossoms for proving correctness of the $O(\sqrt{|V||E|})$ general graph matching algorithms. *Combinatorica*, 14 (1):71–109, 1994.
- [VW12] V. Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In *STOC'12: Proceedings of the 44th Annual ACM Symposium on the Theory of Computing*, 2012.
- [Wri97] S. J. Wright. *Primal-Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, 1997.
- [Ye97] Y. Ye. *Interior Point Algorithms: Theory and Analysis*. John Wiley & Sons, 1997.

A Proof of Lemma 2.1

Let C^* be the value of right-hand side of the equality we need to establish, and - for notational convenience - let us denote the energy $\mathcal{E}_r(\mathbf{f}^*)$ as E^* . So, our goal is to show that $C^* = 1/E^*$ and that taking $\tilde{\phi}$ attains the minimum C^* .

We start by noting that, for any vertex potentials ϕ , we have

$$\sum_{(u,v) \in E} f_{(u,v)}^*(\phi_v - \phi_u) = \sum_v \phi_v \left(\sum_{e \in E^+(v)} f_e^* - \sum_{e \in E^-(v)} f_e^* \right) = \sum_v \phi_v \sigma_v = \boldsymbol{\sigma}^T \boldsymbol{\phi}, \quad (51)$$

where we used the fact that \mathbf{f}^* is a $\boldsymbol{\sigma}$ -flow (cf. (1)).

Note that by the above calculations and the definition of $\tilde{\phi}$ we have

$$\boldsymbol{\sigma}^T \tilde{\phi} = \frac{1}{E^*} \sum_{(u,v) \in E} f_{(u,v)}^* (\phi_v^* - \phi_u^*) = \frac{1}{E^*} \sum_{(u,v) \in E} r_{(u,v)} (f_{(u,v)}^*)^2 = 1, \quad (52)$$

where we used (5) and the definition of energy (4). Therefore, we see that $C^* \leq 1/E^*$ as by (6)

$$\sum_{e=(u,v) \in E} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} = \frac{1}{(E^*)^2} \sum_{e=(u,v) \in E} \frac{(\phi_v^* - \phi_u^*)^2}{r_e} = 1/E^*. \quad (53)$$

Now, let $\hat{\phi}$ be the potential such that $\sum_{(u,v) \in E} \frac{(\hat{\phi}_v - \hat{\phi}_u)^2}{r_{(u,v)}} = C^*$ and let $\hat{\mathbf{f}}$ with $\hat{f}_{(u,v)} := \frac{\hat{\phi}_v - \hat{\phi}_u}{r_{(u,v)}}$, for each $(u,v) \in E$, be the corresponding flow induced via 5. (Note that in principle $\hat{\mathbf{f}}$ does not need to be a $\boldsymbol{\sigma}$ -flow).

From (51) we get that

$$(\mathbf{f}^*)^T \mathbf{R} \hat{\mathbf{f}} = \sum_e r_e f_e^* \hat{f}_e = \sum_{(u,v) \in E} f_{(u,v)}^* (\hat{\phi}_v - \hat{\phi}_u) = \boldsymbol{\sigma}^T \hat{\phi} = 1, \quad (54)$$

where we again used (5) and the fact that $\boldsymbol{\sigma}^T \hat{\phi} = 1$ by definition.

We claim that the energy $\mathcal{E}_r(\hat{\mathbf{f}})$ of $\hat{\mathbf{f}}$ (and thus the value of C^*) is at least $1/E^*$. To this end, let us note that

$$\begin{aligned} \mathcal{E}_r(\hat{\mathbf{f}}) &= \hat{\mathbf{f}}^T \mathbf{R} \hat{\mathbf{f}} = \left(\frac{\mathbf{f}^*}{E^*} + \hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right)^T \mathbf{R} \left(\frac{\mathbf{f}^*}{E^*} + \hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) \\ &= \frac{(\mathbf{f}^*)^T \mathbf{R} \mathbf{f}^*}{(E^*)^2} - 2 \frac{(\mathbf{f}^*)^T}{E^*} \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) + \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right)^T \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right). \end{aligned}$$

As we have seen in (54), $(\mathbf{f}^*)^T \mathbf{R} \hat{\mathbf{f}} = 1$, thus $\frac{(\mathbf{f}^*)^T}{E^*} \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) = 0$ and we can write

$$\begin{aligned} \mathcal{E}_r(\hat{\mathbf{f}}) &= \frac{(\mathbf{f}^*)^T \mathbf{R} \mathbf{f}^*}{(E^*)^2} - 2 \frac{(\mathbf{f}^*)^T}{E^*} \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) + \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right)^T \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) \\ &= \frac{1}{E^*} + \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right)^T \mathbf{R} \left(\hat{\mathbf{f}} - \frac{\mathbf{f}^*}{E^*} \right) \geq \frac{1}{E^*}, \end{aligned}$$

as $\mathbf{f}^T \mathbf{R} \mathbf{f} \geq 0$ for any \mathbf{f} .

So, $C^* \geq 1/E^*$ too and thus $C^* = 1/E^*$. Also, by (52) and (53) we see that $\tilde{\phi}$ indeed attains the minimum, as desired.

B Proof of Corollary 3.4

Let \mathbf{f} be a fractional feasible s - t flow of value F in G and let us consider first the case when F is integral. Recall that the reduction presented in Section 4 allows one to obtain in $\tilde{O}(m)$ time an instance of bipartite \mathbf{b} -matching problem – corresponding to some bipartite graph \bar{G} – that has a property that if there exists a feasible s - t flow of value F in G then \bar{G} has a perfect \mathbf{b} -matching. Now,

the crucial observation is that the proof of that property presented in Section 4 is fully constructive and, in particular, provides an $\tilde{O}(m)$ -time algorithm that produces such a perfect \mathbf{b} -matching in \bar{G} out of a feasible s - t flow in G of value F . Furthermore, this construction also works for fractional flows, it just produces a perfect \mathbf{b} -matching that is fractional.

In the light of the above, we can simply apply this transformation to our flow \mathbf{f} and get a fractional perfect \mathbf{b} -matching \mathbf{x} in \bar{G} . Next, we can use the rounding procedure from Theorem 3.3 to obtain in $\tilde{O}(m)$ time a perfect \mathbf{b} -matching \mathbf{x}^* in \bar{G} that is integral. (Note that since \mathbf{b} is always integral, so is the size of any perfect \mathbf{b} -matching.) This, in turn, allows us to utilize another property of the graph \bar{G} that was established in Section 4. Namely, that out of any integral perfect \mathbf{b} -matching in \bar{G} , one can extract – in $\tilde{O}(m)$ time – an integral and feasible s - t flow \mathbf{f}^* in G of value F . Clearly, by combining all of the above steps, we get our desired integral s - t flow.

Finally, to deal with the case when F is not integral, we just add an arc (s, t) to G , set its capacity to 1, and put a flow of $\lceil F \rceil - F \leq 1$ on it. Obviously, now we have a feasible s - t flow of value $\lceil F \rceil$ in such modified graph G and $\lceil F \rceil$ is integral. Therefore, we can use our approach we described above to get an integral and feasible s - t flow \mathbf{f}^* in this graph and \mathbf{f}^* will have a value of $\lceil F \rceil$. Note that \mathbf{f}^* can have non-zero flow on the arc (s, t) that we added, but as this arc has capacity of 1, there can be exactly one unit of flow on this arc. So, if we simply remove it from \mathbf{f}^* , we will get an integral and feasible s - t flow in the original graph G and the value of \mathbf{f}^* will be $\lceil F \rceil - 1 = \lfloor F \rfloor$, as desired. This concludes the proof of the corollary.

C Appendix to Section 4

C.1 Correctness Analysis

It is easy to verify that the produced \mathbf{b} -matching instance is indeed bipartite (we have edges only between different sides of bipartition P and Q), has exactly $2(m+n-1) = \Theta(m)$ vertices, $3m+n-2 \leq 4m$ edges, and $\|\mathbf{b}\|_1 \leq 4\|\mathbf{u}\|_1$. So, we just need to establish the claimed connection to existence of feasible s - t flows in the graph G .

From Flow \mathbf{f} to Perfect \mathbf{b} -matching \mathbf{x}

To this end, assume that there exists a feasible s - t flow \mathbf{f} in G of value F . To see that a perfect \mathbf{b} -matching in \bar{G} exists, consider a \mathbf{b} -matching \mathbf{x} that, for each arc $e = (u, v)$ in G , takes exactly f_e edges (p_e, q_e) and $u_e - f_e$ edges (q_u, p_e) and (q_e, p_v) . Then, for every vertex v of G other than s and t , \mathbf{x} takes $\sum_{e \in E^+(v)} f_e$ copies of the edge (p_v, q_v) .

To see that \mathbf{x} is indeed a perfect \mathbf{b} -matching, observe that due to feasibility of \mathbf{f} (cf. (2)), $0 \leq f_e \leq u_e$ for each arc e , and thus $\mathbf{x} \geq 0$. Also, by the construction of \mathbf{x} , all vertices p_e and q_e have exactly u_e edges adjacent to them in \mathbf{x} . So, they are fully matched. To see that all vertices p_v and q_v are fully matched too, consider some $v \neq s, t$. Indeed, by definition of \mathbf{x} , we have exactly $\sum_{e \in E^+(v)} u_e - f_e + \sum_{e \in E^+(v)} f_e = \sum_{e \in E^+(v)} u_e = b_{p_v}$ (resp. $\sum_{e \in E^-(v)} u_e - f_e + \sum_{e \in E^-(v)} f_e = \sum_{e \in E^-(v)} u_e = b_{q_v}$) edges adjacent to p_v (resp. q_v), where we used the fact that $\sum_{e \in E^+(v)} f_e = \sum_{e \in E^-(v)} f_e$, as \mathbf{f} obeys flow conservation constraints (1). Finally, in the case of vertex q_s (resp. p_t) we have that their degree in \mathbf{x} is exactly $\sum_{e \in E^-(s)} u_e - f_e = (\sum_{e \in E^-(s)} u_e) - F = b_{q_s}$ (resp. $\sum_{e \in E^+(t)} u_e - f_e = (\sum_{e \in E^+(t)} u_e) - F = b_{p_t}$), due to the value $\sum_{e \in E^-(s)} f_e = \sum_{e \in E^+(t)} f_e$ of the flow \mathbf{f} being exactly F . So, indeed such \mathbf{x} is perfect, as claimed.

From Perfect \mathbf{b} -matching \mathbf{x} to Flow \mathbf{f}

Now, to see that given a perfect \mathbf{b} -matching \mathbf{x} in \bar{G} we can quickly, i.e., in $\tilde{O}(m)$ time, recover an s - t flow of value F that is feasible in G , consider a flow \mathbf{f} given by $f_e = x_{(p_e, q_e)}$ for each arc e in G . That is, the flow f_e on an arc e is equal to the number of times a copy of an edge (p_e, q_e) appears in \mathbf{x} . Note that as the demands b_{p_e} and b_{q_e} of the endpoints of each edge (p_e, q_e) are equal to u_e , \mathbf{f} is feasible in G .

Finally, to prove that \mathbf{f} also preserves flow conservation constraints (cf. (1)), note that as \mathbf{x} is perfect, it has to be that for any vertex v and $e \in E^+(v)$ (resp. $e \in E^-(v)$) $x_{(p_v, q_e)} = b_{q_e} - x_{(p_e, q_e)} = u_e - f_e$ (resp. $x_{(p_e, q_v)} = b_{p_e} - x_{(p_e, q_e)} = u_e - f_e$). So, if we do not take into account the edges (p_v, q_v) , each vertex p_v (resp. q_v) has exactly $\sum_{e \in E^+(v)} b_{q_e} - x_{(p_e, q_e)} = \sum_{e \in E^+(v)} u_e - f_e$ (resp. $\sum_{e \in E^-(v)} b_{p_e} - x_{(p_e, q_e)} = \sum_{e \in E^-(v)} u_e - f_e$) edges adjacent to it in \mathbf{x} . This means, in particular, that in case of q_s (resp. p_t) we need to have that $\sum_{e \in E^-(s)} u_e - f_e = b_{q_s} = (\sum_{e \in E^-(s)} u_e) - F$ (resp. $\sum_{e \in E^+(t)} u_e - f_e = b_{p_t} = (\sum_{e \in E^+(t)} u_e) - F$) and thus $\sum_{e \in E^-(s)} f_e = F$ (resp. $\sum_{e \in E^+(t)} f_e = F$), i.e., the value of \mathbf{f} is F . Furthermore, for any vertex v other than s and t , as \mathbf{x} is perfect, we need to have that $\sum_{e \in E^+(v)} u_e = b_{p_v} = x_{(p_v, q_v)} + \sum_{e \in E^+(v)} u_e - f_e$ (resp. $\sum_{e \in E^-(v)} u_e = b_{q_v} = x_{(p_v, q_v)} + \sum_{e \in E^-(v)} u_e - f_e$). Therefore, $\sum_{e \in E^+(v)} f_e = x_{(p_v, q_v)} = \sum_{e \in E^-(v)} f_e$, i.e., \mathbf{f} obeys all flow conservation constraints. So, indeed \mathbf{f} is a feasible s - t flow of value F in G , as desired.

Lastly, it is worth pointing out that even though in the above proof we assume that both the s - t flow \mathbf{f} and the \mathbf{b} -matching \mathbf{x} are integral, the proof goes through unchanged in the case when \mathbf{f} and \mathbf{x} are fractional. We just will have that if \mathbf{f} is fractional then so will be the corresponding \mathbf{b} -matching \mathbf{x} and vice versa.

D Appendix to Section 5

D.1 Proof of Lemma 5.4

Let us take \mathbf{s}^0 to be the all-ones vector $\mathbf{1}$ (this corresponds to \mathbf{y}^0 assigning zero value to all vertices). Next, let the flow \mathbf{f}^0 and measures ν^0 be defined as follows.

For each arc of the form (s_p, t_q) in \hat{G} , we give it a measure of one in ν^0 and a flow of one unit is sent through it in \mathbf{f}^0 . Now, for each vertex s_p (resp. t_q) in \hat{G} , let $r_p := |\hat{E}^-(s_p)| - 1 - b_p$ (resp. $r_q := |\hat{E}^+(t_q)| - 1 - b_q$). If $r_p \geq 0$ (resp. $r_q \geq 0$) then we put a flow of one and measure of one on the arc (s_p, v^*) (resp. (v^*, t_q)) and a flow and measure of $r_p + 1$ (resp. $r_q + 1$) on the arc (v^*, s_p) (resp. (t_q, v^*)). On the other hand, if $r_p < 0$ (resp. $r_q < 0$) then we put a flow and measure of $1 - r_p$ (resp. $1 - r_q$) on the arc (s_p, v^*) (resp. (v^*, t_q)) and a flow and measure of one on the arc (v^*, s_p) (resp. (t_q, v^*)).

One can verify that the resulting flow \mathbf{f}^0 is indeed a $\hat{\sigma}$ -flow (again, one needs to use here the fact that $\sum_p b_p = \sum_q b_q$, as otherwise there is no perfect \mathbf{b} -matching in G) and thus the solution is primal-dual feasible.

Also, the total measure $\sum_e \nu_e^0$ of all the arcs is at most $\hat{m} + |\mathbf{b}|_1 \leq 3\hat{m}$. Finally, we have that $\hat{\mu}_e^0 = \frac{f_e^0 s_e^0}{\nu_e^0} = 1$ for all arcs e and thus the solution is indeed 0-centered and $\hat{\mu}(\mathbf{f}^0, \mathbf{s}^0, \nu^0) = 1$, as desired.

E Appendix to Section 6

E.1 Proof of Lemma 6.3

That $(1 - \gamma)\alpha \leq \beta \leq (1 + \gamma)\alpha$ follows directly from Fact 5.3.

Let us show now that $\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') = \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$. To this end, let us define γ_e so as

$$(1 + \gamma_e) := \frac{\beta}{\alpha} = \frac{f_e s_e}{\nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}.$$

By definition (10), we have that

$$\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') = \frac{\sum_g f'_g s'_g}{\sum_g \nu'_g} = \frac{(\sum_{g \neq e} f_g s_g) + (1 + \alpha) f_e s_e}{(\sum_{g \neq e} \nu_g) + (1 + \beta) \nu_e} = \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \left(\frac{(\sum_g \nu_g) + \alpha \nu_e (1 + \gamma_e)}{(\sum_g \nu_g) + \alpha \nu_e (1 + \gamma_e)} \right) = \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$$

Now, to bound the centrality of $(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}')$, we need to bound the value of $\|\hat{\boldsymbol{\mu}}' - \hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') \mathbf{1}\|_{\boldsymbol{\nu}', 2}$. However, as $\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') = \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ and the two solutions coincide on all the arcs except e , it suffices to analyze the change in contribution of arc e to the centrality of the solution. Namely, we just need to show that

$$\nu'_e \left(\frac{f'_e s'_e}{\nu'_e} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \right)^2 = (1 + \beta) \nu_e \left(\frac{(1 + \alpha) f_e s_e}{(1 + \beta) \nu_e} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \right)^2 \leq \nu_e \left(\frac{f_e s_e}{\nu_e} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \right)^2.$$

To this end, observe the right side of the above inequality is just

$$\nu_e \left(\frac{f_e s_e}{\nu_e} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \right)^2 = \nu_e \gamma_e^2 \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})^2.$$

So, we need to show that

$$\frac{\nu'_e}{\nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})^2} \left(\frac{f'_e s'_e}{\nu'_e} - \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \right)^2 = (1 + \beta) \left(\frac{(1 + \alpha)(1 + \gamma_e)}{(1 + \beta)} - 1 \right)^2 \leq \gamma_e^2.$$

But this is true as

$$(1 + \beta) \left(\frac{(1 + \alpha)(1 + \gamma_e)}{(1 + \beta)} - 1 \right)^2 = \frac{((1 + \alpha)(1 + \gamma_e) - 1 - (1 + \gamma_e)\alpha)^2}{1 + (1 + \gamma_e)\alpha} = \frac{\gamma_e^2}{1 + (1 + \gamma_e)\alpha} \leq \gamma_e^2,$$

where we used that fact that $(1 + \gamma_e) \geq \frac{1}{2}$ since, by Fact 5.3, $(1 + \gamma_e) \geq (1 - \gamma) \geq \frac{1}{2}$.

E.2 Proof of Lemma 6.5

Note that by construction of the graph \widehat{G} and by Invariant 6.4, we have that for any two vertices v, v' in \widehat{G} there is a directed path from v to v' , as well as, a one from v' and v , with each one of them consisting of at most two arcs and having length at most 4. As $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ is $\boldsymbol{\sigma}$ -feasible then it is, in particular, dual feasible. So, this implies that if \mathbf{y} is the embedding of the vertices of \widehat{G} into a line corresponding to the slack variables \mathbf{s} , then $|y_v - y_{v'}|$ is at most 4 as well. Thus, we can conclude that for any arc $e = (v, v')$ in \widehat{G} , we have that $s_e = \hat{l}_e - y_{v'} + y_v \leq \hat{l}_e + |y_v - y_{v'}|$ is at most 6, as desired.

E.3 Proof of Lemma 6.6

Note first that all the arcs in the original version of \widehat{G} have length 1 and α -stretching can only increase these lengths. This means that it is still true – as in the proof of Lemma 5.1 – that $\hat{\ell}(\mathbf{f}) - \frac{|b|_1}{2}$ is an upper bound on the total flow between the vertices s_p and t_q that is not flowing over the direct arcs (s_p, q_t) reflecting the original edges of G .

Furthermore, as by Invariant 6.4 the total increase in the length of the arcs of \widehat{G} is $\widetilde{O}(\widehat{m}^{\frac{1}{2}-\eta})$, the cost $\hat{\ell}(\mathbf{f}^*)$ of the flow that encodes the perfect \mathbf{b} -matching in G (cf. the proof of Lemma 5.1), can increase to at most $\frac{|b|_1}{2} + \widetilde{O}(\widehat{m}^{\frac{1}{2}-\eta})$. (We use here the fact that f^* never flows more than one unit of flow through any of the arcs.) So, we can assume that the cost $\hat{\ell}(\mathbf{f})$ of the flow \mathbf{f} we have is also $\frac{|b|_1}{2} + \widetilde{O}(\widehat{m}^{\frac{1}{2}-\eta})$, as otherwise we could conclude that no perfect \mathbf{b} -matching exists in G .

However, then it must be the case that the total flow in \mathbf{f} that does not correspond to taking the direct arcs is at most $\widetilde{O}(\widehat{m}^{\frac{1}{2}-\eta})$. Thus, the fractional \mathbf{b} -matching obtained by taking only the flow that uses these direct flow-paths will still result in a near-perfect \mathbf{b} -matching in G .

E.4 Proof of Lemma 6.8

Note first that the upperbound follows directly from Lemma 5.7 as long as we ensure that $C_E > 4$ (which indeed will be the case).

To establish the lowerbound, let us note first that without loss of generality we can assume that in our \mathbf{b} -matching instance $G = (P \cup Q, E)$, $|P| \geq |Q|$ and thus, as our graph G is sparse, $|P| = \Omega(\widehat{m})$. (Otherwise, we just exchange the roles of P and Q in what follows below.)

Let us define a vector of resistances $\tilde{\mathbf{r}}^t$ given by

$$\tilde{r}_e^t := \begin{cases} r_e^t & \text{if } e \in \widehat{E}(s_p), \text{ for some } p \in P \\ 0 & \text{otherwise,} \end{cases}$$

where $\widehat{E}(s_p) := \widehat{E}^+(s_p) \cup \widehat{E}^-(s_p)$ is the set of all arcs incident to s_p in \widehat{G} . In other words, $\tilde{\mathbf{r}}^t$ corresponds to setting to zero resistances (i.e., collapsing) of all the arcs that are not adjacent to some vertex in P ; and making the resistances of arcs that are adjacent to such s_p equal to their original resistances in \mathbf{r}^t . This means, in particular, that $\tilde{r}_e^t \leq r_e^t$, for each arc e , and thus, by Rayleigh Monotonicity principle (cf. Fact 2.2), we know that if $\tilde{\mathbf{f}}^t$ is the electrical $\hat{\sigma}$ -flow determined by resistances $\tilde{\mathbf{r}}^t$ then

$$\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t) \leq \mathcal{E}_{\mathbf{r}^t}(\hat{\mathbf{f}}^t).$$

Therefore, we can just focus on lowerbounding $\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t)$. To this end, note that after collapsing all the arcs that were not adjacent to some vertex in P , we can think of \widehat{G} as a graph that consists only of vertices from P and a single vertex w^* that represents the remaining collapsed vertices. Furthermore, as there is no arcs in \widehat{G} between different vertices s_p , all the arcs in this collapsed graph are of the form (s_p, w^*) or (w^*, s_p) for some $p \in P$.

As a result, $\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t)$ is equal to

$$\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t) = \sum_{p \in P} \sum_{e \in \widehat{E}(s_p)} r_e^t (\tilde{f}_e^t)^2 = \sum_{p \in P} R_p \sigma_{s_p}^2,$$

where R_p is the effective resistance between vertex s_p and w^* with respect to resistances $\tilde{\mathbf{r}}^t$ and the last equality follows as $\tilde{\mathbf{f}}^t$ is a $\hat{\sigma}$ -flow and all arcs are connecting to w^* .

Now, to lowerbound R_p , for some $p \in P$, note that by definition of \mathbf{r}^t (14), the fact that $\nu_e^t \geq 1$ for all e , and Fact 5.3, we have that

$$\frac{1}{R_p} = \sum_{e \in \widehat{E}(s_p)} \frac{1}{r_e^t} = \sum_{e \in \widehat{E}(s_p)} \frac{(f_e^t)^2}{\mu_e^t} \leq \sum_{e \in \widehat{E}(s_p)} \frac{(f_e^t)^2}{(1 - \hat{\gamma}) \nu_e^t \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)} \leq \frac{F_p^2}{(1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)},$$

where $F_p := \sum_{e \in \widehat{E}(s_p)} f_e^t$ and we used the well-known formula for effective resistance of a circuit that consists solely of parallel arcs.

So, all the above considerations allow us to observe that

$$\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t) = \sum_{p \in P} R_p \hat{\sigma}_{s_p}^2 \geq \sum_{p \in P} \frac{(1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)}{F_p^2} \geq (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \frac{|P|^3}{F^2},$$

where $F = \sum_p F_p$ and we used the fact that $|\hat{\sigma}_{s_p}| \geq b_p \geq 1$, as well as, that for any n -dimensional vector \mathbf{x} , $\sum_{i=1}^n \frac{1}{x_i^2} \geq \frac{n^3}{|\mathbf{x}|_1^2}$.

Thus, it remains to provide an upperbound on F . To this end, let us decompose the flow \mathbf{f}^t into flow-paths (whose endpoints are vertices s_p and t_q) and flow-cycles. Clearly, the total contribution of the flow-paths to F can be at most $|\hat{\boldsymbol{\sigma}}|_1 \leq |\mathbf{b}|_1 = O(\widehat{m})$, since our \mathbf{b} -matching instance is balanced. On the other hand, as length of any flow-cycle is at least two and each flow-cycle contributes its whole volume to the duality gap $(\mathbf{f}^t)^T \mathbf{s}^t = \sum_e \mu_e^t$ (as flow-cycles do not exist in optimal solution), the total contribution of flow-cycle to F is at most $\frac{1}{2} \sum_e \mu_e^t$. Thus, by (13) and Invariant 5.2, this contribution is at most $\frac{1}{2} \sum_e \mu_e^t = \frac{1}{2} \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) (\sum_e \nu_e^t) \leq 2\widehat{m} \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \leq 2\widehat{m}$.

Therefore, we can conclude that $F = O(\widehat{m})$ and since $|P|$ is $\Omega(\widehat{m})$ we have

$$\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t) \geq (1 - \hat{\gamma}) \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \frac{|P|^3}{F^2} \geq C_E^{-1} \widehat{m} \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t),$$

where $C_E > 4$ is an appropriately chosen constant.

E.5 Proof of Lemma 6.9

Let us denote by \mathbf{r} , $\boldsymbol{\nu}$, and $\hat{\mathbf{f}}$, respectively, the resistances \mathbf{r}^t , measures $\boldsymbol{\nu}^t$, and electrical $\hat{\boldsymbol{\sigma}}$ -flow $\hat{\mathbf{f}}^t$ before stretch-boost and by \mathbf{r}' , $\boldsymbol{\nu}'$, and $\hat{\mathbf{f}}'$ the corresponding objects after stretch-boost. Also, let as define $S^* := S_{l^*}(\hat{\mathbf{f}}^t) \cap E_H^t$, where l^* is the index of the stretch-boost. In this notation, we want to show that

$$\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}') \geq \left(1 + \frac{\boldsymbol{\nu}(S^*)}{36 \cdot 2^{2l^*}}\right) \mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}}) \geq \left(1 + \frac{\hat{\theta}^2(\boldsymbol{\nu}(S^*))^{\frac{1}{3}}}{36}\right) \mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}}) \geq \left(1 + \frac{\hat{\theta}^2}{36}\right) \mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}}). \quad (55)$$

Clearly, the last inequality follows as $\nu_e \geq 1$ for all arcs e . To see that the second inequality holds, observe that by (22) we have that

$$\boldsymbol{\nu}(S^*) \geq \hat{\theta}^3 2^{3l^*}$$

and thus

$$\frac{\boldsymbol{\nu}(S^*)}{36 \cdot 2^{2l^*}} \geq \frac{(\boldsymbol{\nu}(S^*))^{\frac{1}{3}} (\hat{\theta}^3 2^{3l^*})^{\frac{2}{3}}}{36 \cdot 2^{2l^*}} = \frac{\hat{\theta}^2(\boldsymbol{\nu}(S^*))^{\frac{1}{3}}}{36}.$$

Therefore, once the first inequality in (55) is established, our lemma will follow by choosing $C_S := \frac{1}{36}$.

So, let us proceed to establishing that inequality. By Lemma 2.1, we know that if ϕ^* is the vector of vertex potentials corresponding to the flow $\hat{\mathbf{f}}$ and resistances \mathbf{r} then

$$\frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})} = \sum_{e=(u,v) \in \hat{E}} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e}, \quad (56)$$

where $\tilde{\phi}_v := \phi_v^*/\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})$, for each v , and $\hat{\sigma}^T \tilde{\phi} = 1$.

Now, consider an arc $e \in S^*$. By (5), the definition of the sets $S_{l^*}(\hat{\mathbf{f}}^t)$ (17), and Fact 5.6, we have that

$$\begin{aligned} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} &= \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})^2} r_e \hat{f}_e^2 \geq \frac{(1 - \hat{\gamma}) \nu_e \hat{m} \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu) \rho(\mathbf{f}, \hat{\mathbf{f}})^2}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})^2} \\ &\geq \left(\frac{(1 - \hat{\gamma}) \nu_e \hat{m} \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)}{4 \cdot 2^{2(l^*+1)} \hat{m} \hat{\mu}(\mathbf{f}, \mathbf{s}, \nu)} \right) \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})} \geq \left(\frac{\nu_e}{18 \cdot 2^{2l^*}} \right) \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})}, \end{aligned}$$

where we also used Lemma 5.7. In other words, the contribution of arc e to the sum in (56) constitutes at least $\frac{\nu_e}{18 \cdot 2^{2l^*}}$ -fraction of this sum.

Next, observe that, by definition of 1-stretching, we need to have that the resistance doubles, i.e., $r'_e = 2r_e$, for all the arcs $e \in S^*$, and remains the same for other arcs, i.e., $r'_e = r_e$, for $e \notin S^*$. This means that

$$\begin{aligned} \sum_{e=(u,v) \in \hat{E}} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r'_e} &= \frac{1}{2} \sum_{e=(u,v) \in S^*} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} + \sum_{e=(u,v) \in \hat{E} \setminus S^*} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r'_e} \\ &= \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})} - \frac{1}{2} \sum_{e=(u,v) \in S^*} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} \leq \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})} \left(1 - \frac{\nu(S^*)}{36 \cdot 2^{2l^*}} \right). \end{aligned}$$

But, by Lemma 2.1, we know that the above estimation provides an upper bound on the value of $\frac{1}{\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}')}$, i.e., we have

$$\frac{1}{\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}')} = \min_{\phi | \hat{\sigma}^T \phi = 1} \sum_{e=(u,v) \in \hat{E}} \frac{(\phi_v - \phi_u)^2}{r'_e} \leq \sum_{e=(u,v) \in \hat{E}} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r'_e} \leq \frac{1}{\mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})} \left(1 - \frac{\nu(S^*)}{36 \cdot 2^{2l^*}} \right),$$

where we used the fact that $\hat{\sigma}^T \tilde{\phi} = 1$, by definition of $\tilde{\phi}$. Multiplying both sides by $\frac{\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}') \mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}})}{\left(1 - \frac{\nu(S^*)}{36 \cdot 2^{2l^*}}\right)}$ and noticing that $\frac{1}{(1-x)} \geq (1+x)$ for any $x \geq 0$, gives us the desired inequality in (55).

E.6 Proof of Lemma 6.11

Let us denote the solution $(\mathbf{f}^t, \mathbf{s}^t, \nu^t)$ by $(\mathbf{f}, \mathbf{s}, \nu)$, the associated electrical flow $\hat{\mathbf{f}}^t$ by $\hat{\mathbf{f}}$, and let \mathbf{r} be the resistances \mathbf{r}^t corresponding to this solution (cf. (14)). Also, let $(\mathbf{f}', \mathbf{s}', \nu')$, $\hat{\mathbf{f}}'$, and \mathbf{r}' , denote these respective object after the interior-point method step is applied.

In this notation, our goal is to show that

$$\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}') \geq (1 + C_P \hat{\theta}^2 \ln \hat{m})^{-1} \mathcal{E}_{\mathbf{r}}(\hat{\mathbf{f}}).$$

To perform such lowerbounding of the energy decrease, we proceed similarly as we did in the proof of Lemma 6.9. Namely, by Lemma 2.1, we know that

$$\frac{1}{\mathcal{E}_r(\hat{\mathbf{f}})} = \sum_{e=(u,v) \in \hat{E}} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e},$$

where $\tilde{\phi}_e := \phi_e^*/\mathcal{E}_r(\hat{\mathbf{f}})$ and $\boldsymbol{\sigma}^T \tilde{\boldsymbol{\phi}} = 1$. We want to show that if we keep the same vertex potentials $\tilde{\boldsymbol{\phi}}$ and change the resistances to \mathbf{r}' then still the corresponding sum – as in the equation above – will not increase by too much (and thus provide a good upperbound on $\frac{1}{\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}')}$).

More specifically, recall that by Theorem 5.5, for any arc e ,

$$\frac{r'_e}{r_e} = \frac{(1 + \kappa_e^t)}{(1 - \delta^t)},$$

and that $\|\kappa^t\|_\infty \leq \frac{1}{2}$.

So, by Lemma 2.1, we have that

$$\frac{1}{\mathcal{E}_{\mathbf{r}'}(\hat{\mathbf{f}}')} \leq \sum_{e=(u,v) \in \hat{E}} \frac{(\phi_v - \phi_u)^2}{r'_e} \leq \sum_{e=(u,v) \in \hat{E}} \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} \leq \sum_{e=(u,v) \in \hat{E}} \frac{(1 + 2|\kappa_e^t|)(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e},$$

where we use the fact that $(1 - x)^{-1} \leq (1 + 2x)$ when $x \leq \frac{1}{2}$ and that by definition of $\tilde{\boldsymbol{\phi}}$, $\boldsymbol{\sigma}^T \tilde{\boldsymbol{\phi}} = 1$.

Furthermore, by (5) and definition of $\tilde{\boldsymbol{\phi}}$, we have

$$\sum_{e=(u,v) \in \hat{E}} (1 + 2|\kappa_e^t|) \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} = \frac{1}{\mathcal{E}_r(\hat{\mathbf{f}})} \left(1 + 2 \sum_e |\kappa_e^t| \frac{(\tilde{\phi}_v - \tilde{\phi}_u)^2}{r_e} \mathcal{E}_r(\hat{\mathbf{f}}) \right) = \frac{1}{\mathcal{E}_r(\hat{\mathbf{f}})} \left(1 + 2 \sum_e |\kappa_e^t| \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \right).$$

So, we again just need to show that

$$\sum_e |\kappa_e^t| \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \leq C_P \hat{\theta}^2 \ln \hat{m},$$

and the lemma will follow.

To establish this last claim, note that for any arc e , $\frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})}$ is just the fraction of energy of the flow $\hat{\mathbf{f}}$ (with respect to resistances \mathbf{r}) that is contributed by the arc e . So, by Fact 5.6 and Lemma 6.8, we have that

$$\frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \leq \frac{(1 + \hat{\gamma})\nu_e \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \rho(\mathbf{f}, \hat{\mathbf{f}})_e^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \leq \frac{2C_E \nu_e}{2^{2k}}, \quad (57)$$

whenever $e \in S_k(\hat{\mathbf{f}})$ (cf. (17)), for some integer k .

As a result, we can conclude that

$$\begin{aligned}
\sum_e |\kappa_e^t| \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} &\leq 2 \sum_l \sum_{e \in T_l} \frac{r_e(\hat{f}_e)^2}{2^l \mathcal{E}_r(\hat{\mathbf{f}})} \leq 2 \left(\sum_{l \leq \lfloor \log \hat{\theta}^{-2} \rfloor} \sum_{e \in T_l} \frac{r_e(\hat{f}_e)^2}{2^l \mathcal{E}_r(\hat{\mathbf{f}})} + \sum_{l \geq \lceil \log \hat{\theta}^{-2} \rceil} \sum_{e \in T_l} 2^{\hat{\theta}^2} \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \right) \\
&\leq 2 \left(\sum_{l \leq \lfloor \log \hat{\theta}^{-2} \rfloor} \sum_{e \in T_l} \frac{\nu_e r_e(\hat{f}_e)^2}{2^l \nu_e \mathcal{E}_r(\hat{\mathbf{f}})} \right) + 4\hat{\theta}^2 \leq 4C_E \left(\sum_{l \leq \lfloor \log \hat{\theta}^{-2} \rfloor} \sum_k \sum_{e \in T_l \cap S_k(\hat{\mathbf{f}})} \frac{\nu_e}{2^l 2^{2k}} \right) + 4\hat{\theta}^2 \\
&= 4C_E \left(\sum_{l \leq \lfloor \log \hat{\theta}^{-2} \rfloor} \sum_k \frac{\nu(T_l \cap S_k(\hat{\mathbf{f}}))}{2^{2k+l}} \right) + 4\hat{\theta}^2,
\end{aligned}$$

where T_l denotes $T_l^{\kappa^t}$ (cf. (23), (57), and the fact that

$$\sum_{l \geq \lceil \log \hat{\theta}^{-2} \rceil} \sum_{e \in T_l} \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \leq \sum_e \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} = \frac{\mathcal{E}_r(\hat{\mathbf{f}})}{\mathcal{E}_r(\hat{\mathbf{f}})} = 1.$$

Now, by $2\hat{\theta}$ -smoothness of $\hat{\mathbf{f}}$ (cf. Definition 6.1), we get that

$$\nu(T_l \cap S_k(\hat{\mathbf{f}})) \leq \lfloor \hat{\theta}^3 2^{3(k+1)} \rfloor, \tag{58}$$

for each l and k . Also, the fact that by Lemma 6.10 κ^t is $O(1)$ -restricted implies that, for any fixed l ,

$$\sum_k \nu(T_l \cap S_k(\hat{\mathbf{f}})) = \nu(T_l) \leq O(2^{3l}).$$

Therefore, we can see that for any l , we have

$$\sum_k \frac{\nu(T_l \cap S_k(\hat{\mathbf{f}}))}{2^{2k+l}} \leq \sum_{k=0}^{k'} \frac{\nu(T_l \cap S_k(\hat{\mathbf{f}}))}{2^{2k+l}},$$

for some $k' = l + \log \hat{\theta}^{-1} + O(1)$. Here, we used the fact that by (58), $\nu(T_l \cap S_k(\hat{\mathbf{f}})) = \emptyset$, if $k < 0$ and that the expression we are bounding will be maximized if the set T_l contains as many high-energy arcs as possible. (Note that due to the constraint $\nu(T_l) = O(2^{3l})$ and the bound (58), T_l can then only contain all the arcs in sets $S_k(\hat{\mathbf{f}})$ for all $k \geq 0$ up to k' .) So, we can conclude that

$$\sum_{k=0}^{\infty} \frac{\nu(T_l \cap S_k(\hat{\mathbf{f}}))}{2^{2k+l}} \leq \sum_{k=0}^{k'} \frac{\hat{\theta}^3 2^{3(k+1)}}{2^{2k+l}} = O(\hat{\theta}^3 2^{\frac{k'}{2}-l}) = O(\hat{\theta}^2).$$

To finish our overall bound, we just need to note that by our above derivation, as well as, the fact that $T_l = \emptyset$ if $l \leq 0$ (as $\|\kappa^t\|_{\infty} \leq \frac{1}{2}$),

$$\sum_e |\kappa_e^t| \frac{r_e(\hat{f}_e)^2}{\mathcal{E}_r(\hat{\mathbf{f}})} \leq 4C_E \left(\sum_{l=1}^{\lfloor \log \hat{\theta}^{-2} \rfloor} \sum_k \frac{\nu(T_l \cap S_k(\hat{\mathbf{f}}))}{2^{2k+l}} \right) + 4\hat{\theta}^2 \leq 4C_E \left(O(\log \hat{\theta}^{-2}) O(\hat{\theta}^2) + \hat{\theta}^2 \right) = C_P \hat{\theta}^2 \ln \hat{m},$$

as desired, once $C_P > 1$ is chosen to be large enough.

E.7 Proof of Lemma 6.12

We start by bounding the increase of measure due to freezing. Let us fix some progress step t and some auxiliary arc e that is in T_l for some $l \leq \log \hat{\theta}^{-2}$, where T_l denotes $T_l^{\bar{\kappa}^t}$, as defined in (23). (Note that only arcs in T_l with $l \leq \log \hat{\theta}^{-2}$ can be frozen at step t .)

By Lemma 6.3 and definition of T_l , the increase of measure resulting from $\widehat{m}^{2\eta}|\bar{\kappa}_e^t|$ -stretching e is at most

$$(1 + \hat{\gamma})\widehat{m}^{2\eta}|\bar{\kappa}_e^t|\nu_e^t \leq \frac{2\widehat{m}^{2\eta}}{2^l}.$$

However, by Lemma 6.10, we know that the vector $\bar{\kappa}^t$ is C_R -restricted. Therefore, the total contribution to measure increase of all the frozen arcs in T_l is at most

$$\frac{2\widehat{m}^{2\eta}}{2^l}C_R2^{3l} \leq O(\widehat{m}^{2\eta}2^{2l}) = O(\widehat{m}^{6\eta}),$$

where we used the fact that $l \leq \log \hat{\theta}^{-2}$.

So, as there is at most $\log \hat{\theta}^{-2}$ different sets T_l that contribute in each progress step, and there is at most $\hat{\theta}^{-2} = \widehat{m}^{2\eta}$ progress steps, the overall increase of measure due to freezing is at most $\widetilde{O}(\widehat{m}^{8\eta})$, as required.

Note that once we establish below that all auxiliary have always f_e^t that is within a factor of C_F of F_A , the fact that F_A is much smaller than F_H will imply that all auxiliary arcs are always light and thus never get stretch-boosted. So, the measure of auxiliary arc can increase only due to freezing and we have already bounded this increase above.

Now, to prove the first part of the lemma, let us fix some auxiliary arc e . Initially, f_e^t is equal to F_A . So, one just need to argue that the total multiplicative change of f_e^t during the course of the $\hat{\theta}$ -improvement phase execution is bounded by a constant.

To this end, note that the flows on arcs change only during the progress steps. So, by Theorem 5.5, if we fix some auxiliary arc e , its overall flow changes by a factor of at most

$$\prod_{t=t_0}^{t_f} (1 + |\bar{\kappa}_e^t|).$$

Therefore, the total change of flow of e during progress steps that have not resulted in freezing it, can be bounded by

$$\prod_{t=t_0}^{t_f} (1 + |\bar{\kappa}_e^t|) \leq (1 + \hat{\theta}^2)^{\hat{\theta}^{-2}} \leq \exp(1),$$

that is constant, as desired.

So, now we just need to focus on bounding the change of flow on e resulting from the remaining progress steps, i.e., the ones in which it was frozen. To this end, recall that whenever $|\bar{\kappa}_e^t| \geq \hat{\theta}^{-2}$ in some step t then freezing $\widehat{m}^{2\eta}|\bar{\kappa}_e^t|$ -stretches e . By Lemma 6.3, the resulting increase of measure of e is at least by a factor of

$$(1 + (1 - \hat{\gamma})\widehat{m}^{2\eta}|\bar{\kappa}_e^t|),$$

while the change of the flow is by a factor of at most

$$(1 + |\bar{\kappa}_e^t|).$$

Therefore, as the former factor is significantly larger than the latter one, $\|\bar{\kappa}^t\|_\infty \leq \frac{1}{2}$ (by Theorem 5.5), and as from discussion above we know that once the measure of an arc becomes larger than $C_R \hat{\theta}^{-6}$ it will never be frozen again, the constant bound of the maximum multiplicative change of the flow of an auxiliary arc follows. This concludes the proof of the lemma.

E.8 Proof of Lemma 6.13

By definition of $\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)$ (cf. (10)), we have that

$$\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) = \frac{\sum_e \nu_e^t \hat{\mu}_e^t}{\sum_e \nu_e^t} = \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \left(1 + \frac{\sum_e \nu_e^t \lambda_e}{\sum_e \nu_e^t} \right),$$

where $\lambda_e := \frac{(\hat{\mu}_e^t - \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t))}{\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)}$, for each arc e .

On the other hand, we have that

$$\hat{\mu}(\mathbf{f}', \mathbf{s}', \boldsymbol{\nu}') = \frac{\sum_{e \in S} \nu_e^t \hat{\mu}_e^t}{\sum_{e \in S} \nu_e^t} = \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t) \left(1 + \frac{\sum_{e \in S} \nu_e^t \lambda_e}{\sum_{e \in S} \nu_e^t} \right) \leq \hat{\lambda} \hat{\mu}(\mathbf{f}^{t_0}, \mathbf{s}^{t_0}, \boldsymbol{\nu}^{t_0}) \left(1 + \frac{\sum_{e \in S} \nu_e^t |\lambda_e|}{\boldsymbol{\nu}^t(S)} \right), \quad (59)$$

where S is the set of non-auxiliary arcs of \bar{G} .

Now, observe that by definition of $\hat{\gamma}$ -centrality (cf. (11)) we have

$$\|\boldsymbol{\lambda}\|_{\boldsymbol{\nu}^t, 2} = \frac{\|\hat{\boldsymbol{\mu}}^t - \hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)\|_{\boldsymbol{\nu}^t, 2}}{\hat{\mu}(\mathbf{f}^t, \mathbf{s}^t, \boldsymbol{\nu}^t)} \leq \hat{\gamma}.$$

So, by applying Cauchy-Schwarz inequality we get that

$$\frac{\sum_{e \in S} \nu_e^t |\lambda_e|}{\boldsymbol{\nu}^t(S)} \leq \sqrt{\frac{\sum_{e \in S} \nu_e^t \lambda_e^2}{\boldsymbol{\nu}^t(S)}} = \sqrt{\frac{\|\boldsymbol{\lambda}\|_{\boldsymbol{\nu}^t, 2}^2}{\boldsymbol{\nu}^t(S)}} \leq \sqrt{\frac{\hat{\gamma}^2}{\boldsymbol{\nu}^t(S)}} \leq O(\hat{m}^{-\frac{1}{2}}),$$

where we use the fact that $\boldsymbol{\nu}^t(S) \geq \hat{m}$.

By putting the above inequality and (59) together, the lemma follows.

E.9 Handling Approximate Nature of Electrical Flow Computations

Here, we discuss how one can adjust our algorithm developed in Sections 5–7 to nearly-linear time electrical flow computations that are only approximate – as in Theorem 2.3 – instead of being exact.

To this end, let us first recall that we are using electrical flow computations in two places of our algorithm. One is our improvement step described in Section 7. There, to make the descent step, we compute the electrical flow $\hat{\mathbf{f}}^t$ associated with our solution and then, to make the centering step, we compute the electrical flow $\hat{\mathbf{f}}^t$. The other place where we use electrical flow computations is to check the $\hat{\theta}$ -smoothness condition (cf. Definition 6.1), that is to check which arcs are in the sets $S_l(\hat{\mathbf{f}}^t)$, for $l \leq \log \hat{\theta}^{-3}$. (Note that it is sufficient for us to know this classification only approximately, say up to a constant factor.)

Observe that in all these three cases, we end up computing some electrical $\boldsymbol{\sigma}$ -flows that are determined by some resistances \mathbf{r} defined as $r_e = \frac{s_e}{f_e}$ (cf. (14)), for each arc e , and where $(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})$ is some γ -centered and $\bar{\boldsymbol{\sigma}}'$ -feasible solution with $\gamma \leq \frac{1}{2}$ and both $|\boldsymbol{\sigma}|_1$ and $|\boldsymbol{\sigma}'|_1$ being $O(\hat{m})$. Furthermore, we always have that $\frac{1}{O(\hat{m})} \leq \hat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu}) \leq 1$, all variables s_e are bounded by a constant (see

Lemma 6.5), and the duality gap is $O(\widehat{m})$. (All the definitions that are relevant here can be found in Section 5 and at the beginning of Section 6.)

This implies that, for any arc e , f_e is always polynomially bounded in \widehat{m} . This is so since, given our polynomially-bounded demands, any flow of value $\omega(\widehat{m})$ would need to consist mostly of flow-cycles, and such flow-cycles would contribute to duality gap (as they cannot exist in optimal solution), which is always $O(\widehat{m})$.

This, in turn, together with γ -centrality (see Fact 5.3) and Invariant 5.2, allows us to conclude that all the resistances

$$r_e = \frac{s_e}{f_e} = (1 \pm \gamma) \frac{\nu_e \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}{f_e^2} = (1 \pm \gamma) \frac{s_e^2}{\nu_e \widehat{\mu}(\mathbf{f}, \mathbf{s}, \boldsymbol{\nu})}$$

are within a polynomial in \widehat{m} factor of each other.

It is known (see, e.g., Theorem 2.3 in [CKM⁺11]) that once all the resistances are within polynomial of each other, one can afford very good (and fast) approximation to all the major characteristics of the electrical flows (including good approximation to the flow on each of the edges). In particular, one is able to easily perform (approximate) classification of arcs into sets $S_l(\widehat{\mathbf{f}}^t)$, for $l \leq \log \widehat{\theta}^{-3}$. (Note that we want to classify here only arcs that contribute significant portion of the total energy anyway.) Also, looking at our analysis of our improvement step in Section 7, one can see that the most fundamental requirement there is that the flows $\widehat{\mathbf{f}}^t$ and $\widetilde{\mathbf{f}}^t$ that we compute are indeed electrical flow, i.e., there are voltages that induce them via (5). After all, this is what ensures that our first-order updates to the centrality are canceling out. The fact that these flow might not have the exact demands we requested is of lesser importance. The only effect of the latter will be that our improvement steps will end up perturbing the $\widehat{\boldsymbol{\sigma}}$ -feasibility of our maintained solution. However, given that we have polynomially bounded resistance ratio and logarithmic dependence on error, we can always make these perturbation very small and just fix them at the end of each $\widehat{\theta}$ -improvement steps via the fixing procedure that we already employ to fix the effects of the preconditioning – see Section 6.2.

In the light of the above, we can conclude that indeed, having approximate, instead of exact, electrical flow computations is acceptable for our algorithm, at least as long the dependence of the running time on the error is only logarithmic (which is the case here).

F Appendix to Section 7

To prove the second part of the theorem, note first that indeed $\boldsymbol{\nu}^{t+1} = \widetilde{\boldsymbol{\nu}}^t = \boldsymbol{\nu}^t$. Next, we can check that the cumulative changes of the vectors \mathbf{f}^t and \mathbf{s}^t are equal to

$$\begin{aligned} f_e^{t+1} &= (1 - \delta^t) \left(1 + \frac{\delta^t \widehat{f}_e^t}{(1 - \delta^t) f_e^t} \right) \left(1 - \frac{\widehat{\mu}_e^t - \widehat{\mu}(\widehat{\mathbf{f}}^t, \widehat{\mathbf{s}}^t, \widetilde{\boldsymbol{\nu}}^t)}{\widehat{\mu}_e^t} \right) \left(1 + \frac{\widetilde{f}_e^t}{\widetilde{f}_e^t} \right) f_e^t \\ s_e^{t+1} &= \left(1 - \frac{\delta^t \widehat{f}_e^t}{(1 - \delta^t) f_e^t} \right) \left(1 - \frac{\widetilde{f}_e^t}{\widetilde{f}_e^t} \right) s_e^t, \end{aligned}$$

for each arc e in \widehat{G} .

As a result, by (14), we have that for each arc e ,

$$(1 - \delta^t)r_e^{t+1} = \frac{(1 - \delta^t)s_e^{t+1}}{f_e^{t+1}} = \frac{\left(1 - \frac{\delta^t \hat{f}_e^t}{(1 - \delta^t)f_e^t}\right) \left(1 - \frac{\tilde{f}_e^t}{f_e^t}\right) s_e^t}{\left(1 + \frac{\delta^t \hat{f}_e^t}{(1 - \delta^t)f_e^t}\right) \left(1 - \frac{\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)}{\hat{\mu}_e^t}\right) \left(1 + \frac{\tilde{f}_e^t}{f_e^t}\right) f_e^t}.$$

Recall that from the discussion above we already know that, for each arc e , $\left|\frac{\delta^t \hat{f}_e^t}{(1 - \delta^t)f_e^t}\right| = \frac{\delta^t \rho(\tilde{\mathbf{f}}^t, \mathbf{f}^t)_e}{(1 - \delta^t)} \leq 2\sqrt{\hat{\gamma}} \leq \frac{1}{10}$ (cf. (41)), $\left|\frac{\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)}{\hat{\mu}_e^t}\right| \leq \frac{\hat{\gamma}}{(1 - \hat{\gamma})} \leq \frac{1}{40}$ (cf. Fact 5.3), and $\left|\frac{\tilde{f}_e^t}{f_e^t}\right| = \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e \leq \frac{1}{40}$ (cf. (50)). So, as $r_e^t = \frac{s_e^t}{f_e^t}$, we have that

$$(1 + \kappa_e^t) = \frac{\left(1 - \frac{\delta^t \hat{f}_e^t}{(1 - \delta^t)f_e^t}\right) \left(1 - \frac{\tilde{f}_e^t}{f_e^t}\right)}{\left(1 + \frac{\delta^t \hat{f}_e^t}{(1 - \delta^t)f_e^t}\right) \left(1 - \frac{\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)}{\hat{\mu}_e^t}\right) \left(1 + \frac{\tilde{f}_e^t}{f_e^t}\right)}$$

and by performing a simple Taylor expansion approximation we can obtain that

$$|\kappa_e^t| \leq 2 \left(\frac{\delta^t \rho(\tilde{\mathbf{f}}^t, \mathbf{f}^t)_e}{(1 - \delta^t)} + \frac{|\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)|}{\hat{\mu}_e^t} + \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e \right) \leq 4 \left(\delta^t \rho(\tilde{\mathbf{f}}^t, \mathbf{f}^t)_e + \hat{\kappa}_e^t \right),$$

for each arc e , where $\hat{\kappa}_e^t := \frac{|\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)|}{\hat{\mu}_e^t} + \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e$.

Clearly, this means, in particular, that $\|\kappa_e^t\|_\infty \leq 2\left(\frac{2}{10} + \frac{1}{40} + \frac{1}{40}\right) = \frac{1}{2}$, as desired. So, we just need to show that $\|\hat{\kappa}^t\|_{\nu^t, 2} \leq \frac{1}{16}$ too. To this end, observe that

$$\begin{aligned} \|\hat{\kappa}^t\|_{\nu^t, 2}^2 &\leq 2 \sum_e \nu_e^t \left(\frac{(\hat{\mu}_e^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t))^2}{(\hat{\mu}_e^t)^2} + \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e^2 \right) \\ &\leq 2 \left(\frac{\|\hat{\mu}^t - \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)\mathbf{1}\|_{\nu^t, 2}^2}{(1 - \hat{\gamma})^2 \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)^2} + \frac{1}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)} \sum_e \hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t) \rho(\tilde{\mathbf{f}}^t, \tilde{\mathbf{f}}^t)_e^2 \right) \\ &\leq 2 \left(9\hat{\gamma}^2 + \frac{\mathcal{E}_{\tilde{\mathbf{r}}^t}(\tilde{\mathbf{f}}^t)}{\hat{\mu}(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)} \right) \leq \frac{1}{256}, \end{aligned}$$

as desired, where we used a combination of Fact 5.3, the fact that $(\tilde{\mathbf{f}}^t, \tilde{\mathbf{s}}^t, \tilde{\mathbf{v}}^t)$ is $3\hat{\gamma}$ -centered, as well as, equations (46) and (49). This concludes the proof of the theorem.