



# Optimal Electrical Oblivious Routing on Expanders

Cella Florescu  

ETH Zürich, Switzerland

Rasmus Kyng  

ETH Zürich, Switzerland

Maximilian Probst Gutenberg  

ETH Zürich, Switzerland

Sushant Sachdeva  

University of Toronto, Canada

---

## Abstract

In this paper, we investigate the question of whether the electrical flow routing is a good oblivious routing scheme on an  $m$ -edge graph  $G = (V, E)$  that is a  $\Phi$ -expander, i.e. where  $|\partial S| \geq \Phi \cdot \text{vol}(S)$  for every  $S \subseteq V$ ,  $\text{vol}(S) \leq \text{vol}(V)/2$ . Beyond its simplicity and structural importance, this question is well-motivated by the current state-of-the-art of fast algorithms for  $\ell_\infty$  oblivious routings that reduce to the expander-case which is in turn solved by electrical flow routing.

Our main result proves that the electrical routing is an  $O(\Phi^{-1} \log m)$ -competitive oblivious routing in the  $\ell_1$ - and  $\ell_\infty$ -norms. We further observe that the oblivious routing is  $O(\log^2 m)$ -competitive in the  $\ell_2$ -norm and, in fact,  $O(\log m)$ -competitive if  $\ell_2$ -localization is  $O(\log m)$  which is widely believed.

Using these three upper bounds, we can smoothly interpolate to obtain upper bounds for every  $p \in [2, \infty]$  and  $q$  given by  $1/p + 1/q = 1$ . Assuming  $\ell_2$ -localization in  $O(\log m)$ , we obtain that in  $\ell_p$  and  $\ell_q$ , the electrical oblivious routing is  $O(\Phi^{-(1-2/p)} \log m)$  competitive. Using the currently known result for  $\ell_2$ -localization, this ratio deteriorates by at most a sublogarithmic factor for every  $p, q \neq 2$ .

We complement our upper bounds with lower bounds that show that the electrical routing for any such  $p$  and  $q$  is  $\Omega(\Phi^{-(1-2/p)} \log m)$ -competitive. This renders our results in  $\ell_1$  and  $\ell_\infty$  unconditionally tight up to constants, and the result in any  $\ell_p$ - and  $\ell_q$ -norm to be tight in case of  $\ell_2$ -localization in  $O(\log m)$ .

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Routing and network design problems

**Keywords and phrases** Expanders, Oblivious routing for  $\ell_p$ , Electrical flow routing

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2024.65

**Category** Track A: Algorithms, Complexity and Games

**Related Version** *Full Version*: <https://arxiv.org/abs/2406.07252>

**Funding** *Rasmus Kyng*: “Algorithms and complexity for high-accuracy flows and convex optimization” grant (no. 200021 204787) of the Swiss National Science Foundation.

*Maximilian Probst Gutenberg*: “Algorithms and complexity for high-accuracy flows and convex optimization” grant (no. 200021 204787) of the Swiss National Science Foundation.

*Sushant Sachdeva*: NSERC Discovery Grant RGPIN-2018-06398, an Ontario Early Researcher Award (ERA) ER21-16-284, and a Sloan Research Fellowship.

**Acknowledgements** We would like to thank Yang P. Liu for pointing us to the Riesz-Thorin theorem.

## 1 Introduction

In this paper, we study flow-routing problems on connected, undirected (multi-)graphs  $G = (V, E)$ . A broad and well-studied class of single-commodity flow problems arises by seeking a flow  $f \in \mathbb{R}^E$  that routes given demands  $\chi \in \mathbb{R}^V$ , while minimizing a  $\ell_p$ -norm of



© Cella Florescu, Rasmus Kyng, Maximilian Probst Gutenberg, and Sushant Sachdeva; licensed under Creative Commons License CC-BY 4.0

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).

Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;

Article No. 65; pp. 65:1–65:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



## 65:2 Optimal Electrical Oblivious Routing on Expanders

the flow. Denoting the graph edge-vertex incidence matrix by  $\mathbf{B} \in \mathbb{R}^{V \times E}$ , we can write these optimization problems as

$$\min_{\mathbf{B}\mathbf{f}=\boldsymbol{\chi}} \|\mathbf{f}\|_p. \quad (1)$$

The case  $p = \infty$  is known as undirected maximum flow, while  $p = 2$  is called electrical flow and  $p = 1$  is called transshipment. Here we focus for simplicity on the unweighted setting, but all results in this paper and in related work can in fact be extended to work in weighted graphs.

We can generalize these flow problems to the multi-commodity case by allowing a collection of demands  $\{\boldsymbol{\chi}_i\}$  to be routed simultaneously by a collection of flows  $\{\mathbf{f}_i\}$ , while minimizing a single objective on all of them.

$$\min_{\mathbf{B}\mathbf{f}_i=\boldsymbol{\chi}_i, \forall i} \left\| \sum_i \mathbf{f}_i \right\|_p. \quad (2)$$

For any  $p$ , solutions with  $(1 + 1/\text{poly}(|E|))$ -multiplicative error to these problems can be computed in polynomial time and for the single-commodity setting even in almost-linear time [18, 7]. For the special cases of  $p = 1, 2, \infty$ , optimal solutions can be computed in polynomial time via linear/convex programming.

However, in many settings, we may want to sacrifice optimality of our routing solutions for simplicity of the routing algorithm. A particularly simple and popular approach is *oblivious routing*, where a collection of routing paths are chosen in advance between every pair of nodes, without knowing the demands that will be eventually routed. Historically, oblivious routings were first studied on specific networks, specifically the hypercube [37, 38]. A deeply influential technique in this area is the work of Racké [27]. An *oblivious routing* is linear operator  $\mathbf{A} \in \mathbb{R}^{E \times V}$  that maps any valid<sup>1</sup> demand vector  $\boldsymbol{\chi} \in \mathbb{R}^V$  to a flow  $\mathbf{f} = \mathbf{A}\boldsymbol{\chi}$  that routes  $\boldsymbol{\chi}$ . This extends to routing multiple demands in the multi-commodity setting,  $\{\mathbf{f}_i = \mathbf{A}\boldsymbol{\chi}_i\}$ .

Conceptually, a highly attractive feature is that multiple demand vectors can be routed simultaneously without knowing the other demands, and a single demand can be broken down into multiple terms, e.g. source-sink demand pairs, and routings of each pair can be again computed separately. These features make oblivious routings ideal for online routing problems – which was the original motivation for their construction [27].

As mentioned above, using oblivious routing comes at the sacrifice of optimality. To get a quantitative measure of the loss a routing scheme  $\mathbf{A}$  might incur in the  $\ell_p$  metric, we define the *competitive ratio* of  $\mathbf{A}$ , denoted by  $\rho_p(\mathbf{A})$ , to be the maximal ratio between the objective value achieved by an oblivious routing  $\mathbf{A}$  and the optimal solution achieved by *any* (multi-commodity) demand.

In a ground-breaking sequence of papers [27, 3, 9], Racké, Azar, Cohen, Fiat, Kaplan, and Englert showed that for all  $\ell_p$ -norms, oblivious routings with competitive ratio  $\tilde{O}(1)$  exist<sup>2</sup>. In fact, for the well-studied setting of  $p = \infty$ , [28] gave an optimal construction with  $O(\log m)$  competitive ratio in polynomial time, matching a  $\Omega(\log m)$  lower bound [4, 25].

### Fast Algorithms and Applications for $\ell_\infty$ Oblivious Routing

$\ell_\infty$  oblivious routings are a fundamental tool in obtaining fast approximate maximum flow algorithms in undirected graphs. Building on the techniques in [35, 24], [16, 33] give algorithms that show that  $O(\text{poly}(\alpha/\varepsilon))$  applications of a  $\alpha$ -competitive  $\ell_\infty$  oblivious routing

<sup>1</sup> A demand  $\boldsymbol{\chi}$  can be routed on a connected graph iff  $\sum_v \chi(v) = 0$ .

<sup>2</sup> We use  $\tilde{O}(\cdot)$  to hide polylogarithmic factors in the graph size  $m$ .

yield  $(1 + \varepsilon)$ -approximate maximum flow on undirected graphs by using gradient descent. They then developed almost-linear time algorithms for  $m^{o(1)}$ -competitive  $\ell_\infty$  oblivious routing. As a result, they obtained approximate undirected maximum flow in time  $m^{1+o(1)} \text{poly}(1/\varepsilon)$  – one of the major recent breakthroughs in modern graph algorithms.

Later, [29] gave a reduction from computing  $\tilde{O}(1)$ -competitive  $\ell_\infty$  oblivious routings to approximate maximum flows resulting in a  $m^{1+o(1)}$  time algorithm. [26] then showed that combining these approaches recursively yields a  $\tilde{O}(m)$  algorithm to compute  $\tilde{O}(1)$ -competitive  $\ell_\infty$  oblivious routings and a  $\tilde{O}(m \text{poly}(1/\varepsilon))$  algorithm for  $(1 + \varepsilon)$ -approximate undirected maximum flow [26]. Recently, [12] presented an alternative, simple  $\tilde{O}(m)$  algorithm to obtain (and maintain)  $\ell_\infty$  oblivious routings with subpolynomial competitive factor.

While recently the first *exact* maximum flow algorithm with runtime  $m^{1+o(1)}$  was given in [7],  $\ell_\infty$  oblivious routings and approximate undirected maximum flow remain important tools with many algorithms crucially relying on them as subroutines to obtain runtime  $\tilde{O}(m)$ .

We point out that above, for simplicity, we did not properly distinguish between  $\ell_\infty$  oblivious routings which are only constructed in [16], and their weaker counterparts *congestion approximators* which are used in all other constructions. A congestion approximator is a linear operator  $C$  that maps each demand  $\chi$  to vector  $c = C\chi$  such that  $\|c\|_\infty$  approximates the objective value of (2). Note, that  $c$  is not necessarily a flow.

### $\ell_\infty$ Oblivious Routing on Expanders and in General Graphs

*Valiant's trick* [38], a popular scheme that routes demands from each source to a set of randomly chosen intermediate nodes before routing them to the destination, establishes the existence of  $O(\Phi^{-1} \log n)$ -competitive  $\ell_\infty$ -oblivious routings in expanders. However, implementing Valiant's trick algorithmically requires computing multi-commodity flows, which are expensive to compute.

To the best of our knowledge, the only fast algorithm that computes an  $\ell_\infty$  oblivious routing on general graphs is given in [16]. In their approach, they first reduce the problem to finding an  $\ell_\infty$  oblivious routing on a  $\Phi$ -expander with unit-weights (in this case  $\Phi = m^{-o(1)}$ ). They then exploit a simple but striking statement, previously demonstrated by Kelner and Maymounkov [15]: the electrical flow routing, henceforth denoted by  $A_\varepsilon$ , on a  $\Phi$ -expander is a  $O(\Phi^{-2} \log m)$ -competitive  $\ell_\infty$ -routing. It was later observed by Schild-Rao-Srivastava [32] that on unweighted graphs, the statement can be derived from Cheeger's Inequality ([5, 1]). Further, the electrical flow routing can be applied efficiently after  $\tilde{O}(m)$  preprocessing, due to the breakthrough result by Spielman and Teng [35] and subsequent work [17, 8, 20, 14]<sup>3</sup>. [16] then demonstrates that by assembling and combining these routings on expanders, one obtains an  $\ell_\infty$  oblivious routing of the entire graph that can be evaluated efficiently.

As far as we know, no other fast algorithm is currently known to compute  $\ell_\infty$  oblivious routing, and all fast algorithms that compute congestion approximators again reduce to expanders on which cuts can be approximated by stars. Therefore, to the best of our knowledge, every almost-linear time approach to constructing  $\ell_\infty$  oblivious routing reduces to expanders, and on expanders the only known fast algorithm for obtaining an  $\ell_\infty$  oblivious routing is to use the electrical flow routing.

<sup>3</sup> Technically, only a high-accuracy solution is obtained which suffices for our application.

### Oblivious Routing for any $\ell_p$

Analogous to the reduction of solving approximate undirected maximum flow via few applications of  $\ell_\infty$  oblivious routing, Sherman later showed in [34], that any  $\ell_p$ -norm minimizing flow on undirected graphs can be computed to an  $(1 + \varepsilon)$ -approximation by applying  $\ell_p$  oblivious routings with  $\alpha$  competitive ratio  $\tilde{O}(\text{poly}(\alpha/\varepsilon))$  times via gradient descent.

While we are not aware of any article studying fast algorithms for the general  $\ell_p$ -norm, the  $\ell_1$ -norm has received considerable attention and  $\tilde{O}(m)$  time algorithms were given with competitive ratio  $\tilde{O}(1)$  [23, 39, 30], and adapted to fully-dynamic graphs in [6].

### Oblivious Routing for $\ell_1$ on Expanders

Further, at least in the unit-capacity setting, the result by Kelner and Maymounkov [15] extends seamlessly to the  $\ell_1$ -norm, i.e. the electric flow routing  $\mathbf{A}_\mathcal{E}$  has competitive ratio  $O(\Phi^{-2} \log m)$ . This follows since the electrical flow routing is given by  $\mathbf{A}_\mathcal{E} = \mathbf{B}^\top \mathbf{L}^+$ , where  $\mathbf{B}$  is the vertex-edge incidence matrix and  $\mathbf{L} = \mathbf{B}\mathbf{B}^\top$  is the Laplacian matrix of the graph and then bounding the  $\ell_\infty$  competitive ratio of the oblivious routing for multicommodity flow problems is equivalent to bounding the quantity  $\|\mathbf{A}_\mathcal{E}\mathbf{B}\|_{\infty \rightarrow \infty}$ , where  $|\cdot|$  denotes the entrywise absolute value, while the  $\ell_1$  competitive ratio equals  $\|\mathbf{A}_\mathcal{E}\mathbf{B}\|_{1 \rightarrow 1}$ . The matrix  $\Pi = \mathbf{A}_\mathcal{E}\mathbf{B} = \mathbf{B}^\top \mathbf{L}^+ \mathbf{B}$  is a frequently-studied orthogonal projection matrix and it is a symmetric matrix, since  $\mathbf{B}^\top \mathbf{L}^+ \mathbf{B} = \mathbf{B}^\top \mathbf{A}_\mathcal{E}^\top = (\mathbf{A}_\mathcal{E}\mathbf{B})^\top = (\mathbf{B}^\top \mathbf{L}^+ \mathbf{B})^\top$  where we use that  $\mathbf{L}^+$  is symmetric. But it is further well-known that  $\|\mathbf{X}\|_{\infty \rightarrow \infty} = \|\mathbf{X}^\top\|_{1 \rightarrow 1}$  and thus we have that  $\|\Pi\|_{\infty \rightarrow \infty} = \|\Pi\|_{1 \rightarrow 1}$ , i.e. the competitive ratios achieved by  $\mathbf{A}_\mathcal{E}$  in  $\ell_1$ - and  $\ell_\infty$ -norm are equal.

### Beyond Oblivious Routing

The quantity  $\|\Pi\|_{\infty \rightarrow 1}$  is important in several other contexts: It captures the so-called *localization* of electrical flow on the graph [32]. Localization measures the  $\ell_1$ -length of the electrical flow corresponding to a demand placed at two endpoints of an edge, averaged over all edges. The bound  $\rho_1(\mathbf{A}_\mathcal{E}) = O(\Phi^{-2} \log m)$  implies a stronger statement in expanders: For every such edge-demand, the  $\ell_1$ -length is bounded by  $O(\Phi^{-2} \log m)$ . It is known that in general graphs, localization is bounded by  $O(\log^2 m)$  – but it is open whether  $O(\log m)$  holds (a lower bound of  $\Omega(\log m)$  is known) although it is widely believed. From [15], we see that any graph with expansion  $1/o(\sqrt{\log m})$  achieves localization  $o(\log^2 m)$ .

Localization has been used in a number of contexts, including sampling random spanning trees in almost-linear time [31], computing spectral subspace sparsification [22] of Laplacian matrices, and building oblivious routings using  $\tilde{O}(\sqrt{m})$  electrical flows [11].

An interesting message of our paper is that electrical flow on expanders simultaneously is an excellent  $\ell_1$  and  $\ell_\infty$  oblivious routing, i.e. it uses flow paths are both short and low congestion. A broad theory of expanders that simultaneously allow for short and low-congestion paths has recently been developed in [10, 13], allowing for other possible trade-offs between length and congestion than those obtained by conventional expanders.

## 1.1 Main Contributions

In this article, we study a simple but important question:

*Given any  $p \in [1, \infty]$ , what is the competitive ratio  $\rho_p(\mathbf{A}_\mathcal{E})$  of the electrical flow routing  $\mathbf{A}_\mathcal{E}$  on a  $\Phi$ -expander?*

We first settle this question for the important cases when  $p \in \{1, \infty\}$  by proving the following theorem that proves an upper bound that is tight up to constant factors.

► **Theorem 1.1.** *For a  $\Phi$ -expander multigraph  $G = (V, E)$  with edge-vertex incidence matrix  $\mathbf{B}$  and Laplacian  $\mathbf{L}$ , the electrical routing  $\mathbf{A}_{\mathcal{E}} = \mathbf{B}^{\top} \mathbf{L}^+ \mathbf{B}$  has competitive ratios  $\rho_{\infty}$  and  $\rho_1$  for multi-commodity  $\ell_{\infty}$  and  $\ell_1$  routing both bounded by*

$$\rho_{\infty}(\mathbf{A}_{\mathcal{E}}), \rho_1(\mathbf{A}_{\mathcal{E}}) \leq 3 \cdot \frac{\log(2m)}{\Phi}$$

The Riesz-Thorin theorem then gives us a way to smoothly interpolate between the upper bounds of any two  $\ell_{p_1}$ - and  $\ell_{p_2}$ -norm competitive ratios to obtain an upper bound on the  $\ell_p$ -norm competitive ratio  $\rho_p(\mathbf{A}_{\mathcal{E}})$  for any  $p_1 < p < p_2$ . Using a smooth interpolation between our results for  $\ell_1$ - and  $\ell_{\infty}$ -norm, we thus obtain the following more general result.

► **Theorem 1.2.** *For a  $\Phi$ -expander multigraph  $G = (V, E)$ , and any  $p \in [1, \infty]$ , we have that the competitive ratio of  $\mathbf{A}_{\mathcal{E}}$  is*

$$\rho_p(\mathbf{A}_{\mathcal{E}}) \leq 3 \cdot \frac{\log(2m)}{\Phi}.$$

It was proven in [32], that  $\Pi = \mathbf{B}^{\top} \mathbf{L}^+ \mathbf{B}$  satisfies  $\|\Pi\|_{2 \rightarrow 2} \leq O(\log^2 n)$  (where  $|\cdot|$  indicates entry-wise absolute value). We refer to this quantity  $\|\Pi\|_{2 \rightarrow 2}$  as  $\ell_2$ -localization. It is widely believed that  $\|\Pi\|_{2 \rightarrow 2} \leq O(\log m)$ . Implicit in earlier works, albeit perhaps not widely observed, is that  $\rho_2(\mathbf{A}_{\mathcal{E}}) = \|\Pi\|_{2 \rightarrow 2}$  (see Lemma 4.1), i.e. the competitive ratio of multi-commodity  $\ell_2$  routing is exactly characterized by  $\ell_2$ -localization. By interpolating with this norm bound and our bounds on  $\rho_1(\mathbf{A}_{\mathcal{E}})$  and  $\rho_{\infty}(\mathbf{A}_{\mathcal{E}})$ , we obtain potentially much stronger bounds on the competitive ratio  $\rho_p(\mathbf{A}_{\mathcal{E}})$ .

► **Corollary 1.3** (implied by Riesz-Thorin). *For a  $\Phi$ -expander multigraph  $G = (V, E)$ , and any  $p \in [2, \infty]$  and  $q$  given by  $1/p + 1/q = 1$ , we have that the competitive ratios of  $\mathbf{A}_{\mathcal{E}}$  for the  $\ell_p$  and  $\ell_q$  norms are*

$$\rho_p(\mathbf{A}_{\mathcal{E}}), \rho_q(\mathbf{A}_{\mathcal{E}}) \leq \|\Pi\|_{2 \rightarrow 2}^{2/p} \left( 3 \cdot \frac{\log(2m)}{\Phi} \right)^{1-2/p}.$$


We complement our upper bounds with strong, unconditional lower bounds. Remarkably, if  $\|\Pi\|_{2 \rightarrow 2} \leq O(\log m)$ , as widely believed, then our bounds are tight up to constants. Even with the currently known fact  $\|\Pi\|_{2 \rightarrow 2} \leq O(\log^2 n)$ , our lower bounds still prove a sublogarithmic gap in every constant  $p \neq 2$  and  $q$  and optimal  $\Phi$  dependency.

► **Theorem 1.4.** *For an infinite number of positive integers  $n$  and any  $\Phi \in [1/\sqrt[3]{n}, 1]$ , for any  $p \in [2, \infty]$  and  $q$  given by  $1/p + 1/q = 1$ , we have that*

$$\rho_p(\mathbf{A}_{\mathcal{E}}), \rho_q(\mathbf{A}_{\mathcal{E}}) \geq \Omega \left( \frac{\log m}{\Phi^{1-2/p}} \right).$$

## 1.2 Oblivious Electric Routing in Weighted Graphs

We can extend Theorem 1.1 to weighted graphs in the following way: Consider a graph  $G = (V, E)$  with positive integer edge capacities  $\mathbf{c} \in \mathbb{R}^E$  and positive integer edge lengths  $\mathbf{s} \in \mathbb{R}^E$ . Letting  $\mathbf{C}, \mathbf{S} \in \mathbb{R}^{E \times E}$  denote diagonal matrices with  $\mathbf{c}$  and  $\mathbf{s}$  on the diagonal respectively, we are interested in the optimal weighted  $\ell_{\infty}$ - and  $\ell_1$ -routings for given demands  $D = \{\chi_1, \dots, \chi_k\}$



■ **Figure 1** An illustration of the result given in Corollary 1.3 of the different competitive ratios achieved with respect to each  $\ell_p$ -norm, where  $n$  and  $\Phi$  are fixed and  $\Phi \ll 1/\log m$ . The red curve shows the optimal ratios, achieved if localization is in  $O(\log m)$ , which is also obtained up to constant factors by our lower bound in Theorem 1.4. The lilac curve shows the current trade-off where we use the known result that localization is in  $O(\log^2 n)$ . For  $\Phi \gg 1/\log m$  the upper bound of Theorem 1.2 achieves values between the two curves.

$$\text{opt}_\infty(D) = \min_{\mathbf{B}\mathbf{f}_i = \mathbf{x}_i, \forall i} \|\mathbf{C}^{-1} \sum_i \mathbf{f}_i\|_\infty \text{ and } \text{opt}_1(D) = \min_{\mathbf{B}\mathbf{f}_i = \mathbf{x}_i, \forall i} \|\mathbf{S} \sum_i \mathbf{f}_i\|_1 \quad (3)$$

Now consider defining an electrical routing by choosing resistances  $\mathbf{R} = \mathbf{S}\mathbf{C}^{-1}$  and defining the electrical routing  $\mathbf{A}_\mathcal{E} = \mathbf{R}^{-1}\mathbf{B}^\top(\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\top)^+$ . Let  $\hat{G}$  denote the multigraph with edge  $e$  replaced by  $\mathbf{c}(e)$  unit-weight paths of length  $\mathbf{s}(e)$ . Now, one can easily show that the electrical routing in  $G$  according to  $\mathbf{A}_\mathcal{E}$  is equal to the unit-weight electrical routing in  $\hat{G}$ , when mapping flows on a capacitated edge to a collection of flows on multi-edge paths.

▶ **Corollary 1.5** ((Informal) Electrical Oblivious  $\ell_\infty$ - and  $\ell_1$ -Routing on Weighted Expanders). For (multi-)graph  $G = (V, E)$  with edge-vertex incidence matrix  $\mathbf{B}$  and positive integer edge weights and lengths given as diagonal matrices  $\mathbf{C}, \mathbf{S} \in \mathbb{R}^{E \times E}$ , the electrical routing  $\mathbf{A}_\mathcal{E} = \mathbf{R}^{-1}\mathbf{B}^\top(\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\top)^+$ , where  $\mathbf{R} = \mathbf{S}\mathbf{C}^{-1}$ , has competitive ratios  $\rho_\infty$  and  $\rho_1$  for multicommodity  $\ell_\infty$  and  $\ell_1$  routing both bounded by

$$\rho_\infty(\mathbf{A}_\mathcal{E}), \rho_1(\mathbf{A}_\mathcal{E}) \leq 3 \cdot \frac{\ln(2|E_{\hat{G}}|)}{\Phi(\hat{G})}$$

where  $\hat{G}$  denotes the multigraph with edge  $e$  replaced by a path of length  $\mathbf{S}(e, e)$  with  $\mathbf{C}(e, e)$  unit-weight multi-edges across each hop of the path.

▶ **Remark 1.6.** When we take all edge lengths  $\mathbf{S}(e, e)$  to be 1, the expansion of  $\hat{G}$  equals the usual definition of expansion in graph  $G$  with edge weight equal to capacity.

### 1.3 New Implications for Localization

From the connection to localization outlined earlier, we also immediately conclude that localization of graphs with  $1/o(\log m)$  expansion improve over the general localization bound of  $O(\log^2 m)$ .

▶ **Corollary 1.7** ((Informal) Localization of Electrical Flow). For a multigraph  $G = (V, E)$ , the average over multi-edges of the  $\ell_1$ -norm the electrical flow routing 1 unit of flow across the multi-edge is bounded by  $O(\min\{\Phi^{-1} \log m, \log^2 n\})$ , and hence graphs with expansion  $\Phi = 1/o(\log m)$  have localization  $o(\log^2 m)$ .

## 1.4 Roadmap

We next give a Preliminary section to set up the necessary notation for the article. We then prove Theorem 3 in Section 3. We use this result together with the Riesz-Thorin theorem to obtain Theorem 1.2 and Corollary 1.3 in Section 4. Finally, in Section 5, we give our lower bounds as stated in Theorem 1.4.

## 2 Preliminaries

### General Definitions

For any  $n \in \mathbb{N}^*$ , we let  $[n]$  denote the set  $\{1, 2, \dots, n\}$ . We let  $\mathbf{1}$  denote the all ones vector and  $\mathbf{1}_S$  denote the vector that has ones in the positions indexed by the elements of the set  $S$  and zeros otherwise. For any  $\mathbf{A} \in \mathbb{R}^{m \times n}$ , we let  $|\mathbf{A}|$  denote the matrix where the absolute value operator has been applied entrywise.

### Graphs

Although our results in the contribution section are for unweighted graphs, we also prove stronger statements in the article that also work on weighted graphs. Therefore, we define various notions with respect to weighted graphs.

Given an input graph  $G = (V, E, w)$  with positive weights which we all assume to be at least 1, we define  $n = |V|$  and  $m = |E|$ . We assume an arbitrary underlying direction assigned to each edge of  $G$ . We define the edge-vertex incidence matrix  $\mathbf{B} \in \mathbb{R}^{V \times E}$  of  $G$  as

$$\mathbf{B}(w, e) = \begin{cases} -1, & \text{if } e = (w, v) \\ 1, & \text{if } e = (v, w) \\ 0, & \text{otherwise} \end{cases}.$$

We define the Laplacian  $\mathbf{L} = \mathbf{B}\mathbf{W}\mathbf{B}^\top$  where  $\mathbf{W}$  is the diagonal matrix given by the weights  $w$  and denote by  $\mathbf{L}^+$  the pseudo-inverse of the Laplacian. We call  $\Pi = \mathbf{B}^\top \mathbf{L}^+ \mathbf{B}$  the *unweighted* projection matrix of  $G$ .

### Expanders

We say  $G$  is a  $\Phi$ -expander if  $|\partial S| \geq \Phi \cdot \text{vol}(S)$  for every  $S \subseteq V$ ,  $\text{vol}(S) \leq \text{vol}(V)/2$ , where we define  $|\partial S|$  to be the weight of all edges with exactly one endpoint in  $S$ , and  $\text{vol}(S)$  to be the sum of weighted degrees of vertices in  $S$ .

### Flows and Congestion

We say  $\chi \in \mathbb{R}^V$  is a demand vector if  $\mathbf{1}^\top \chi = 0$ . We let  $\chi_{(a,b)} \in \mathbb{R}^V$  for every  $(a, b) \in E$  to be the unitary demand on the edge  $(a, b)$ , that is  $\chi_{(a,b)} = \mathbf{1}_a - \mathbf{1}_b$ . We say a vector  $\mathbf{f} \in \mathbb{R}^E$  is a flow that routes demand  $\chi$  if  $\mathbf{B}\mathbf{f} = \chi$ . Given an arbitrary norm  $\|\cdot\|$  on  $\mathbb{R}^E$ , we define the congestion of a multi-set of flows  $\{\mathbf{f}_1, \dots, \mathbf{f}_k\}$  to be:

$$\text{cong}(\{\mathbf{f}_1, \dots, \mathbf{f}_k\}) = \left\| \mathbf{W}^{-1} \sum_{i=1}^k \mathbf{f}_i \right\|.$$

### Oblivious Routings

We define an oblivious routing on  $G$  to be a linear operator  $\mathbf{A} \in \mathbb{R}^{E \times V}$  such that  $\mathbf{BA}\chi = \chi$  for all demand vectors  $\chi \in \mathbb{R}^V$ , i.e. to be a flow that routes the demand  $\chi$ .

Given a multiset of demands  $D = \{\chi_1, \dots, \chi_k\}$ , we define the optimal congestion achievable by

$$\text{opt}(D) = \min_{\{\mathbf{f}_i\}_{i \in [k]} \text{ multiset} : \mathbf{B}\mathbf{f}_i = \chi_i, \forall i} \text{cong}(\{\mathbf{f}_i\}_{i \in [k]}).$$

This allows us to define the competitive ratio of an oblivious routing, which we define

$$\rho(\mathbf{A}) = \max_{\{\chi_i\}_{i \in [k]} \text{ multiset} : \chi_i \perp \mathbf{1}, \forall i} \frac{\text{cong}(\{\mathbf{A}\chi_i\}_{i \in [k]})}{\text{opt}(\{\chi_i\}_{i \in [k]})}.$$

Note that whenever we use the subscript “ $p$ ” for the competitive ratio  $\rho$ , we mean that the norm used in defining the congestion in that special case is the  $\ell_p$ -norm.

### Electrical Flows and Voltages

In this article, we define the electric flow routing operator  $\mathbf{A}_\mathcal{E} = \mathbf{WB}^\top \mathbf{L}^+$ . Right-applying the operator  $\mathbf{A}$  to any demand  $\chi$  yields the electric flow  $\mathbf{f} = \mathbf{A}\chi$  that routes the demand  $\chi$ . We define the electrical energy associated with the flow vector  $\mathbf{f}$  by  $\mathcal{E}(\mathbf{f}) = \mathbf{f}^\top \mathbf{W}^{-1} \mathbf{f}$ .

We define the electric voltage vector  $\mathbf{v} \in \mathbb{R}^V$  with respect to a demand  $\chi$  by  $\mathbf{v} = \mathbf{L}^+ \chi$ . We define the electrical energy associated with the voltage vector  $\mathbf{v}$  as  $\mathcal{E}(\mathbf{v}) = \mathbf{v}^\top \mathbf{L} \mathbf{v}$ . Note in particular that the energy of voltages induced by a certain demand coincides with the energy of the respective flow.

We introduce the notion of “fractional” volume at given a threshold  $t \in \mathbb{R}$  with respect to a given voltage vector  $\mathbf{v} \in \mathbb{R}^V$ . We first define the fractional volume per edge and then for the whole graph. For an edge  $(a, b) \in E$ :

$$\text{vol}_{\geq t}(a, b) = \begin{cases} 2 \cdot w(a, b), & \text{if } \mathbf{v}(a) > t \\ 2 \cdot w(a, b) \cdot \frac{\mathbf{v}(b) - t}{\mathbf{v}(b) - \mathbf{v}(a)}, & \text{if } \mathbf{v}(b) \leq t \\ 0, & \text{otherwise} \end{cases}$$

For the whole graph  $G$ :

$$\begin{aligned} \text{vol}_{\geq t}(t) &= \sum_{(a,b) \in E} \text{vol}_{\geq t}(a, b), \\ \text{vol}_{\geq t}^+(t) &= \text{vol}_{\geq t}(t) + 1. \end{aligned}$$

Note that in our notation we omit specifying which voltage vector the “fractional” volume function is tied to, as it will be clearly specified upon usage during the proofs.

We define  $S_t = \{a \in V \mid \mathbf{v}(a) \geq t\}$  as the set of vertices whose voltages are greater or equal to the arbitrary threshold  $t \in \mathbb{R}$  and the cut determined by the voltage threshold  $t \in \mathbb{R}$  to be  $C_t = (S_t, V \setminus S_t)$ . For convenience, we let the weight of the cut  $C_t$  be  $\delta(t) = |\partial C_t| = \sum_{e \in C_t} w(e)$ .

## 3 An Upper Bound on the Competitive Ratio of Electrical Flow Routing for $\ell_\infty$

In this section, we prove our main technical result, Theorem 1.1, by establishing a tight upper bound on the competitive ratio when the congestion is defined in terms of  $\ell_\infty$  which then by the symmetry of  $\Pi$  immediately gives the same competitive ratio in  $\ell_1$ .



However, while Theorem 1.1 only claims a result for unweighted (multi-)graphs, we show in this section that  $\mathcal{A}_E$  even has good competitive ratio in weighted graphs with respect to  $\ell_\infty$ . However, in the weighted setting, we cannot use the bound for  $\ell_\infty$  to derive a bound on  $\ell_1$ , as the matrix-norms that exactly characterize the competitive ratio are not equal in general. Nonetheless, the “multi-graph” trick from Corollary 1.5 can be used to transform the weighted setting into an unweighted instance and derive bounds.

### Intuition for our Proof

Kelner and Maymoukov showed that in order to bound the congestion of the electrical routing, it suffices, via a duality (or transposition) argument, to bound the worst case  $\ell_1$ -norm of the flow induced by routing 1 unit of flow electrically across any edge. We adopt the same approach, but give a more precise analysis.

Suppose  $e = (x, y)$  is the edge such that routing one unit of flow between the endpoints causes the highest overall congestion. We let  $\mathbf{v}$  be the associated voltage vector that induces the electrical flow routing one unit from  $x$  to  $y$ . The overall congestion then equals  $\sum_{(a,b) \in E} w(a,b)|\mathbf{v}(a) - \mathbf{v}(b)|$ . We can express this by integrating with respect to voltage along a voltage threshold cut with respect to  $\mathbf{v}$ , where the function being integrated at point  $t$  is exactly  $Y_t = \sum_{(a,b) \in C_t} w(a,b)$ , where  $C_t$  is the cut at voltage threshold  $t$ . This ensures that after integrating  $Y_t$  over the entire voltage range, each edge  $(a, b)$  contributes exactly  $w(a,b)|\mathbf{v}(a) - \mathbf{v}(b)|$ , as desired. Our proof proceeds by leveraging that the flow crossing the cut  $C_t$  at threshold  $t$  is exactly  $\sum_{(a,b) \in C_t} w(a,b)|\mathbf{v}(a) - \mathbf{v}(b)|$ .

As we are sending one unit of flow from  $x$  to  $y$ , and all electrical flow goes one way across a voltage cut, this quantity is exactly 1. At each threshold  $t$ , this creates an “on average” relationship between voltage difference  $|\mathbf{v}(a) - \mathbf{v}(b)|$  and weight  $w(a,b)$  for edges being cut. This in turn allows us to establish a pointwise relationship at each threshold voltage  $t$  between the growth in congestion and the change in volume at  $t$ . Armed with this relationship, we can bound the accumulated congestion of the integrated cuts in terms of the accumulated volume, and this yields our result.

### Contrast with the Kelner-Maymoukov Proof

It is instructive to consider why the Kelner-Maymoukov congestion bound loses an additional factor  $\Phi$  compared to our bound. For concreteness, consider the graph given by a direct edge from  $x$  to  $y$  and an additional  $k$  disjoint paths of length  $k$  from  $x$  to  $y$ . It can be shown that in this example, the edge that governs the congestion bound in the strategy above is in fact the direct  $(x, y)$  edge.

Kelner-Maymoukov upper bound the true competitive ratio  $\rho' = \sum_{(a,b) \in E} w(a,b)|\mathbf{v}(a) - \mathbf{v}(b)|$  by the quantity  $\rho'' = \sum_{a \in V} \mathbf{d}(a)|\mathbf{v}(a) - c|$  for some constant  $c$  (see Equation (4.3) in [15]). On this concrete graph,  $\rho'$  can be explicitly evaluated and is  $\Theta(k)$ . As the graph has expansion  $1/k$ , we can think of this as a bound of  $\Theta(1/\Phi)$ . But,  $\rho''$  is  $\Theta(k^2)$  i.e.  $\Theta(1/\Phi^2)$ .

However, Kelner and Maymoukov’s strategy makes it difficult to directly bound  $\rho'$  as they first measure changes in volume over a (discrete) sequence of threshold cuts, and then changes in voltage over the same sequence of cuts. Their discrete sequence of cuts skips entirely over some edges, i.e. there will be edges that are not crossing any of their cuts. This makes it difficult to establish an estimate for each edge of the pointwise relation between its contribution to volume growth versus voltage growth or congestion growth. Hence, they work with summed bounds on voltage and compare these with summed bounds on volume, which naturally yields bounds on  $\rho''$ . But, as we have seen, a bound on  $\rho''$  must inherently be loose as there is a gap between  $\rho'$  and  $\rho''$ .

## 65:10 Optimal Electrical Oblivious Routing on Expanders

► **Theorem 3.1** ( $\ell_\infty$  Competitive bound of electrical flows). *For a weighted  $\Phi$ -expander multigraph  $G = (V, E, w)$ , the following holds:*

$$\rho_\infty(\mathbf{A}_E) \leq 3 \cdot \frac{\ln(\text{vol}(V))}{\Phi}.$$

**Proof.** We first use that

$$\rho_\infty(\mathbf{A}_E) = \max_{e \in E} \left\| \mathbf{W} \mathbf{B}^\top \mathbf{L}^+ \boldsymbol{\chi}_e \right\|_1. \quad (4)$$

as shown in [16] as part of the proof of Lemma 26 (by using primarily Lemmas 10 and 11). In order to prove the desired inequality, we then fix an  $e \in E$  such that the quantity in (4) gets maximized, and let  $\mathbf{v} = \mathbf{L}^+ \boldsymbol{\chi}_e$  be the voltage induced by setting a unitary demand on this edge. Thus, we equivalently aim to bound:

$$\rho_\infty(\mathbf{A}_E) = \left\| \mathbf{W} \mathbf{B}^\top \mathbf{L}^+ \boldsymbol{\chi}_e \right\|_1 = \left\| \mathbf{W} \mathbf{B}^\top \mathbf{v} \right\|_1 = \sum_{(a,b) \in E} w(a,b) \cdot |\mathbf{v}(a) - \mathbf{v}(b)|. \quad (5)$$

Observe now that shifting all the values of  $\mathbf{v}$  by the same constant does not change the value expressed in (5), and therefore we can assume without loss of generality that the voltages are centered around 0, that is we can assume  $\text{vol}(\{i \in V \mid \mathbf{v}(i) \geq 0\}) \geq \text{vol}(V)/2$  and  $\text{vol}(\{i \in V \mid \mathbf{v}(i) \leq 0\}) \geq \text{vol}(V)/2$ .

Note that, by convention, the electrical flow  $\mathbf{f}_E = \mathbf{A}_E \boldsymbol{\chi}_e$  induces an orientation on the edges in the set  $E$ . Henceforth, we assume without loss of generality that orientations of edges in  $E$  align with the direction of the electrical flow  $\mathbf{f}_E$ , that is  $\mathbf{f}_E(a,b) = w(a,b) \cdot (\mathbf{v}(b) - \mathbf{v}(a)) \geq 0$  for any  $(a,b) \in E$ .

In the following, we will employ the definitions introduced in Section 2. These concepts give rise to the notion of “fractional” volume, which will ultimately allow us to bound the quantity of Equation (5).

It can easily be proven that  $\text{vol}_\geq$  is continuous in  $\mathbb{R}$  and differentiable at any threshold level  $t \in \mathbb{R}$  for which there does not exist a node  $a \in V$  such that  $\mathbf{v}(a) = t$ . Furthermore, if  $t_{\min} = \min\{\mathbf{v}(a) \mid a \in V\}$  and  $t_{\max} = \max\{\mathbf{v}(a) \mid a \in V\}$ , then  $\text{vol}_\geq(t_{\min}) = \text{vol}(G)$  and  $\text{vol}_\geq(t_{\max}) = 0$ . We can even assume without loss of generality a more precise centering of the voltages around 0, namely that  $\text{vol}_\geq(0) = \text{vol}(V)/2$ .

A voltage threshold level  $t \in \mathbb{R}$  can be seen as naturally determining a cut  $C_t$  in  $G$ . Note that the assumption about centering the voltages around 0 ensures that  $\text{vol}(S_t) \leq \text{vol}(V)/2$  for any  $t > 0$ , so it holds that  $\min\{\text{vol}(S_t), \text{vol}(V \setminus S_t)\} = \text{vol}(S_t)$ .

By taking the orientation of the edges into account, we can drop the absolute value operator and rewrite Equation (5) as:

$$\begin{aligned} \sum_{(a,b) \in E} w(a,b) \cdot |\mathbf{v}(a) - \mathbf{v}(b)| &= \sum_{(a,b) \in E} w(a,b) \cdot (\mathbf{v}(b) - \mathbf{v}(a)) \\ &= \sum_{(a,b) \in E} w(a,b) \cdot \int_{t_{\min}}^{t_{\max}} \mathbb{1}_{\mathbf{v}(a) < t \leq \mathbf{v}(b)} dt \\ &= \int_{t_{\min}}^{t_{\max}} \sum_{(a,b) \in E} w(a,b) \cdot \mathbb{1}_{\mathbf{v}(a) < t \leq \mathbf{v}(b)} dt \\ &= \int_{t_{\min}}^{t_{\max}} \delta(t) dt \\ &= \int_{t_{\min}}^0 \delta(t) dt + \int_0^{t_{\max}} \delta(t) dt. \end{aligned} \quad (6)$$

We will bound the quantity in (6) by separately bounding each of the two terms in the last equality. Only the proof for the integral over the non-negative values of  $t$  will be presented, the one for the non-positive values proceeds in an analogous manner. Assume thus for the rest of the proof that  $t \geq 0$  holds.

In order to obtain the bound on  $\delta(t)$  for non-negative  $t$ , we will inspect the rate of change of the “fractional” volume with respect to the voltage threshold of the fixed voltage vector  $\mathbf{v}$ . Intuitively, we grow an electrical threshold ball, and directly relate the change in volume to the stretch accumulated at the current voltage threshold. In more precise terms, we compute a bound on the derivative of  $\text{vol}_{\geq}^+$  with respect to  $t$  on the domain of differentiability as follows:

$$\begin{aligned}
-\frac{d}{dt}\text{vol}_{\geq}^+(t) &= -\frac{d}{dt}\text{vol}_{\geq}(t) \\
&= -\frac{d}{dt}\left(\sum_{(a,b) \in E} \text{vol}_{\geq t}(a,b)\right) \\
&= -\frac{d}{dt}\left(\sum_{(a,b) \in C_t} \text{vol}_{\geq t}(a,b)\right) \\
&= \sum_{(a,b) \in C_t} -\frac{d}{dt}\text{vol}_{\geq t}(a,b) \\
&= \sum_{(a,b) \in C_t} 2 \frac{w(a,b)}{\mathbf{v}(b) - \mathbf{v}(a)}.
\end{aligned} \tag{7}$$

By the construction based on the voltage levels, all edges of the cut  $C_t$  have their head in the set  $S_t$ . Therefore, the flow carried by these edges has to be the unit flow, since this is the demand of  $\chi_e$ :

$$1 = \sum_{(a,b) \in C_t} f_{\mathcal{E}}(a,b) = \sum_{(a,b) \in C_t} w(a,b) \cdot (\mathbf{v}(b) - \mathbf{v}(a)). \tag{8}$$

Hereafter, we show that the negative volume change must exceed the square of the cut size. Informally, the change in volume per edge is relatively large whenever the voltage gap across the edge is small (volume change scales as inversely proportional to the gap compared to the cut-value of the edge). But, a “typical” edge in the cut must have a fairly small voltage gap, as we otherwise route too much flow across the gap. Formally, since the voltage drops among the edges in the cut  $C_t$  are non-negative, we can use (8) and the definition of the conductance of  $G$  to further bound (7) using the Cauchy–Bunyakovsky–Schwarz inequality:

65:12 Optimal Electrical Oblivious Routing on Expanders

$$\begin{aligned}
-\frac{d}{dt}\text{vol}_{\geq}^+(t) &= 2 \sum_{(a,b) \in C_t} \frac{w(a,b)}{\mathbf{v}(b) - \mathbf{v}(a)} \\
&= 2 \left( \sum_{(a,b) \in C_t} \frac{w(a,b)}{\mathbf{v}(b) - \mathbf{v}(a)} \right) \cdot 1 \\
&= 2 \left( \sum_{(a,b) \in C_t} \frac{w(a,b)}{\mathbf{v}(b) - \mathbf{v}(a)} \right) \left( \sum_{(a,b) \in C_t} w(a,b) \cdot (\mathbf{v}(b) - \mathbf{v}(a)) \right) \\
&\geq 2 \left( \sum_{(a,b) \in C_t} \sqrt{\frac{w(a,b)}{\mathbf{v}(b) - \mathbf{v}(a)} \cdot w(a,b) \cdot (\mathbf{v}(b) - \mathbf{v}(a))} \right)^2 \\
&\geq 2 \left( \sum_{(a,b) \in C_t} w(a,b) \right)^2 \\
&= 2 \cdot \delta(t)^2 \\
&\geq 2 \cdot \delta(t) \cdot \Phi \cdot \text{vol}(S_t).
\end{aligned} \tag{9}$$

Denote by  $\text{vol}_{\text{int}}(t) = \text{vol}(S_t) - \delta(t)$  twice the weight of the edges that have both endpoints in the set  $S_t$ . Recall that we assumed all of the edges to have weights at least 1, therefore it holds that  $\delta(t) \geq 1$ . The definition of “fractional” volume implies  $\text{vol}_{\geq}(t) \leq \text{vol}_{\text{int}} + 2\delta(t)$ , which can be used to further bound (9):

$$\begin{aligned}
-\frac{d}{dt}\text{vol}_{\geq}^+(t) &\geq 2 \cdot \delta(t) \cdot \Phi \cdot \text{vol}(S_t) \\
&= 2 \cdot \delta(t) \cdot \Phi \cdot (\text{vol}_{\text{int}}(t) + \delta(t)) \\
&= \delta(t) \cdot \frac{2\Phi}{3} \cdot (3\text{vol}_{\text{int}}(t) + 3\delta(t)) \\
&\geq \delta(t) \cdot \frac{2\Phi}{3} \cdot (\text{vol}_{\text{int}}(t) + 2\delta(t) + 1) \\
&\geq \delta(t) \cdot \frac{2\Phi}{3} \cdot (\text{vol}_{\geq}(t) + 1) \\
&= \delta(t) \cdot \frac{2\Phi}{3} \cdot \text{vol}_{\geq}^+(t).
\end{aligned}$$

Observe that we can rewrite the inequality above to obtain a bound on  $\delta(t)$ :

$$\delta(t) \leq -\frac{3}{2\Phi} \cdot \frac{1}{\text{vol}_{\geq}^+(t)} \cdot \frac{d}{dt}\text{vol}_{\geq}^+(t). \tag{10}$$

We can now use equation (10) to bound the integral over the interval  $[0, t_{\max}]$  in (6):

$$\int_0^{t_{\max}} \delta(t) dt \leq \int_0^{t_{\max}} -\frac{3}{2\Phi} \cdot \frac{1}{\text{vol}_{\geq}^+(t)} \cdot \frac{d}{dt}\text{vol}_{\geq}^+(t) dt.$$

The latter integral can easily be computed via the change of variable  $u = \text{vol}_{\geq}^+(t)$ , yielding the integration bounds  $\text{vol}_{\geq}^+(t_{\max}) = \text{vol}_{\geq}(t_{\max}) + 1 = 1$  and  $\text{vol}_{\geq}^+(0) = \text{vol}_{\geq}(0) + 1 = \text{vol}(V)/2 + 1$ :

$$\begin{aligned}
\int_0^{t_{\max}} \delta(t) dt &\leq \int_0^{t_{\max}} -\frac{3}{2\Phi} \cdot \frac{1}{\text{vol}_{\geq}^+(t)} \cdot \frac{d}{dt} \text{vol}_{\geq}^+(t) dt \\
&= \int_{t_{\max}}^0 \frac{3}{2\Phi} \cdot \frac{1}{\text{vol}_{\geq}^+(t)} \cdot \frac{d}{dt} \text{vol}_{\geq}^+(t) dt \\
&= \frac{3}{2\Phi} \int_1^{\text{vol}(V)/2+1} \frac{1}{u} du.
\end{aligned}$$

Since the graph has by assumption at least two nodes connected by an edge with weight at least 1, it follows that  $\text{vol}(V) \geq 2 \iff \text{vol}(V) \geq \text{vol}(V)/2 + 1$ . Coupled with the fact that  $u > 0$  for  $u \in [1, \text{vol}(V)]$ , this gives us the final bound for the integral:

$$\begin{aligned}
\int_0^{t_{\max}} \delta(t) dt &\leq \frac{3}{2\Phi} \int_1^{\text{vol}(V)/2+1} \frac{1}{u} du \\
&\leq \frac{3}{2\Phi} \int_1^{\text{vol}(V)} \frac{1}{u} du \\
&= \frac{3}{2\Phi} \cdot \ln(\text{vol}(V)).
\end{aligned}$$

As already mentioned, the same bound can be obtained for the other term in (6) in an analogous manner.

Combining the aforementioned result with the relations given by (5) and (6) gives the desired inequality:

$$\rho_{\infty}(\mathbf{A}_{\mathcal{E}}) = \int_{t_{\min}}^0 \delta(t) dt + \int_0^{t_{\max}} \delta(t) dt \leq 2 \cdot \frac{3 \ln(\text{vol}(V))}{2\Phi} = \frac{3 \ln(\text{vol}(V))}{\Phi}. \quad \blacktriangleleft$$

#### 4 An Upper Bound on the Competitive Ratio of Electrical Flow Routing for $\ell_p$ (for any $p$ )

In this section, we prove two generalizations of Theorem 3.1. Previously, we gave a bound on the competitive ratio when the congestion was defined in terms of the  $\ell_{\infty}$ -norm. This result can be extended to  $\ell_p$ -norms for an arbitrary  $p \in [1, \infty]$  by using Theorem 3.1, and instantiating a special case of the Riesz-Thorin theorem.

This establishes both the results in Theorem 1.2 and in Corollary 1.3. We stress that the results obtained in this section crucially exploit the symmetry of  $\Pi$  and therefore only hold for unweighted graphs.

##### A Toolbox for $\ell_p$ -Norms

We first use the following lemma. Its proof is fairly standard and follows the structure of the proof of Lemma 24 in [16]. We include a full proof in an extended version of this paper on arXiv.

► **Lemma 4.1** (Competitive ratio of  $\ell_p$ -norms). *Let  $G = (V, E)$  be a multigraph. For any  $p \in [1, \infty]$  and oblivious routing  $\mathbf{A}$ , we have*

$$\rho_p(\mathbf{A}) = \|\|\mathbf{AB}\|\|_{p \rightarrow p}.$$

Our second tool is the Riesz-Thorin theorem. We explicitly state the two relevant special cases of the theorem that we require in the next section for the convenience of the reader.

## 65:14 Optimal Electrical Oblivious Routing on Expanders

► **Theorem 4.2** (Special cases of the Riesz-Thorin theorem, see [36, Theorem 1.3]). *Let  $\mathbf{A} \in \mathbb{R}^{m \times n}$  be a matrix with non-negative entries. For any  $p \in (1, \infty)$  it holds that:*

$$\|\mathbf{A}\|_{p \rightarrow p} \leq \|\mathbf{A}\|_{1 \rightarrow 1}^{\frac{1}{p}} \cdot \|\mathbf{A}\|_{\infty \rightarrow \infty}^{1 - \frac{1}{p}}.$$

Furthermore, for  $p \in (2, \infty)$ ,

$$\|\mathbf{A}\|_{p \rightarrow p} \leq \|\mathbf{A}\|_{2 \rightarrow 2}^{\frac{2}{p}} \cdot \|\mathbf{A}\|_{\infty \rightarrow \infty}^{1 - \frac{2}{p}}.$$

### Proof of Theorem 1.2

We have that the theorem already holds for  $\ell_1$  and  $\ell_\infty$  since we have proven Theorem 1.1 in the previous section. Consider therefore any  $p \in (1, \infty)$ . Then, we have

$$\begin{aligned} \rho_p(\mathbf{A}_\mathcal{E}) &\stackrel{\text{Lemma 4.1}}{=} \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{p \rightarrow p} \\ &\stackrel{\text{Theorem 4.2}}{\leq} \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{1 \rightarrow 1}^{\frac{1}{p}} \cdot \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{\infty \rightarrow \infty}^{1 - \frac{1}{p}} \\ &\stackrel{\text{Lemma 4.1}}{=} (\rho_1(\mathbf{A}_\mathcal{E}))^{\frac{1}{p}} \cdot (\rho_\infty(\mathbf{A}_\mathcal{E}))^{1 - \frac{1}{p}} \\ &\stackrel{\text{Theorem 1.1}}{\leq} \left(3 \cdot \frac{\ln(2m)}{\Phi}\right)^{\frac{1}{p}} \cdot \left(3 \cdot \frac{\ln(2m)}{\Phi}\right)^{1 - \frac{1}{p}} \\ &= 3 \cdot \frac{\ln(2m)}{\Phi}. \end{aligned}$$

### Proof of Corollary 1.3

Consider next any  $p \in (1, \infty)$ . Since we have again that for  $q$  given by  $1/p + 1/q = 1$ , we have  $\|\mathbf{X}\|_{p \rightarrow p} = \|\mathbf{X}^\top\|_{q \rightarrow q}$ , we can assume w.l.o.g. that  $p \geq 2$ . Similarly to [21], we obtain

$$\begin{aligned} \rho_p(\mathbf{A}_\mathcal{E}) &\stackrel{\text{Lemma 4.1}}{=} \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{p \rightarrow p} \\ &\stackrel{\text{Theorem 4.2}}{\leq} \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{2 \rightarrow 2}^{\frac{2}{p}} \cdot \|\mathbf{A}_\mathcal{E} \mathbf{B}\|_{\infty \rightarrow \infty}^{1 - \frac{2}{p}} \\ &\stackrel{\text{Lemma 4.1}}{=} (\rho_2(\mathbf{A}_\mathcal{E}))^{\frac{2}{p}} \cdot (\rho_\infty(\mathbf{A}_\mathcal{E}))^{1 - \frac{2}{p}} \\ &\stackrel{\text{Lemma 4.1, Theorem 1.1}}{\leq} (\|\Pi\|_{2 \rightarrow 2})^{\frac{2}{p}} \cdot \left(3 \cdot \frac{\ln(2m)}{\Phi}\right)^{1 - \frac{2}{p}}. \end{aligned}$$

## 5 A Lower Bound for Competitive Ratio of Electric Flow Routing

Finally, in this section, we provide a strong lower bound on the competitive ratio of the electrical routing scheme in any  $\ell_p$ -norm.

► **Theorem 5.1** (Restatement of Theorem 1.4). *For an infinite number of positive integers  $n$  and any  $\Phi \in [1/\sqrt[3]{n}, 1]$ , for any  $p \in [2, \infty]$  and  $q$  given by  $1/p + 1/q = 1$ , we have that*

$$\rho_p(\mathbf{A}_\mathcal{E}), \rho_q(\mathbf{A}_\mathcal{E}) \geq \Omega\left(\frac{\log m}{\Phi^{1-2/p}}\right).$$

In our proof, we use the following theorem given in [2]. We remind the reader that the *girth* of a graph  $G$  is the weight of the smallest weight cycle of  $G$ .

► **Theorem 5.2** ([2, Theorem 1.2]). *There are infinitely many positive integer  $\Delta$  and  $n$ , for which an  $n$ -vertex unweighted graph  $G_{\Delta, n} = (V, E)$  exists such that  $G$  is  $\Phi_{\text{const}}$ -expander that is  $\Delta$ -regular with  $\Phi_{\text{const}} = \Theta(1)$  such that  $G$  has girth  $\Omega(\log_\Delta n)$ .*

In our proof, we use the existential result behind the statement to refine a proof technique previously used by Englert and Räcke [9] to give a lower bound on the competitive ratio of any  $\ell_p$  oblivious routing scheme. Our refinement can also be used to strengthen their result by a  $\Theta(\log \log n)$  factor.

In our proof, we crucially exploit the following facts about effective resistance. Recall that the effective resistance of a graph  $G$  for a pair  $(s, t) \in V^2$  is the minimum energy required to route one unit of demand from  $s$  to  $t$  in  $G$ , or alternatively the difference in voltages of  $s$  and  $t$  induced by routing this unit of demand via an electrical flow which is given by  $\chi_{(s,t)} \mathbf{L}^+ \chi_{(s,t)}$ . The facts below can be derived straightforwardly from Cheeger's Inequality, mixing of random walks, and characterization of effective resistance by commute times (see for example [19]).

► **Fact 5.3.** *For  $G$  being a constant-degree  $\Omega(1)$ -expander, we have that the effective resistance of any pair  $(s, t) \in V$  is in  $\Theta(1)$ .*

► **Fact 5.4.** *Given two constant-degree graphs  $G$  and  $H$  over the same vertex set  $V$ . If the effective resistance for a pair  $(s, t) \in V^2$  is in  $\Theta(1)$  in both  $G$  and  $H$ , then the electrical flow routing one unit of demand from  $s$  to  $t$  on the union of graphs  $G \cup H$  sends at least a constant fraction of the flow over  $G$  and a constant fraction of the flow over  $H$ .*

Let us now give a lower bound for any  $p \geq 2$  and any parameter  $\Phi \in [1/n, 1]$  such that  $1/\Phi$  is integer. We start by considering the electrical routing  $\mathbf{A}_\mathcal{E}$  for a large constant  $\Delta$  and any  $n$  and  $G_{\Delta,n}$  of the multi-commodity demand that is given by routing for each edge  $e = (u, v)$  in  $G_{\Delta,n}$  one unit of a commodity from  $u$  to  $v$ , i.e. we consider the demand  $\chi = \left\{ \chi_{(u,v)} \right\}_{e=(u,v) \in E(G_{\Delta,n})}$ . Towards understanding the electrical routing, we prove the following simple claim.

▷ **Claim 5.5.** *For any edge  $e$  in  $G_{\Delta,n}$  where  $\Delta$  is a large constant, we have that the electrical flow  $\mathbf{f} = \mathbf{A}_\mathcal{E} \chi_{(u,v)}$  routing the demand  $\chi_{(u,v)}$  has  $\|\mathbf{f}\|_1 = \Omega(\log n)$ .*

*Proof.* The claim follows from showing that  $\mathbf{f}(e)$  carries only  $(1 - \varepsilon)$  units of flow for some constant  $\varepsilon > 0$ . This is because it implies that a constant fraction of the flow is not routed via the edge  $e$ . But since each path between the endpoints of  $e$  that does not use the edge  $e$  is of length  $\Omega(\log n)$  (by the girth bound in Theorem 5.2), we have that this  $\varepsilon$ -fraction adds  $\Omega(\varepsilon \log n) = \Omega(\log n)$  units of flow to the network  $G_{\Delta,n}$ .

To prove the claim, it suffices to observe that the graph  $G_{\Delta,n} \setminus \{e\}$  is a  $\Omega(1)$ -expander. But to this end, it suffices to observe that since the conductance of  $G_{\Delta,n}$  does not depend on  $\Delta$  by Fact 5.3, by choosing  $\Delta$  sufficiently large (i.e. at least twice the inverse of the conductance), we have that each cut contains at least 2 edges and thus the conductance of  $G_{\Delta,n} \setminus \{e\}$  is at least half of the conductance of  $G_{\Delta,n}$ , and thus still constant.

Using that the trivial graph consisting only of the edge  $e$  is a constant-degree  $\Omega(1)$ -expander, we thus have that the effective resistance of the pair  $(u, v)$  in both the graph  $G_{\Delta,n} \setminus e$  and  $e$  is constant by Fact 5.3. Thus, by Fact 5.4, we have that a constant fraction of the demand  $\chi_{(u,v)}$  is not routed into  $e$ , as desired. ◁

Using that multi-commodity flows do not cancel, we thus have that each edge in  $G_{\Delta,n}$  carries on average  $\Omega(\log n)$  units of flow. We next transform the graph  $G_{\Delta,n}$  to then obtain our final gadget on which we can prove the lower bound.

► **Definition 5.6.** *Let  $G_{\Delta,n}^\Phi$  be the graph obtained from  $G_{\Delta,n}$  by replacing each edge with  $1/\Phi$  vertex-disjoint paths of length  $1/\Phi$  between the endpoints of the vertices. Thus,  $G_{\Delta,n}^\Phi$ , for  $\Delta$  being a constant, has  $\Theta(n/\Phi^2)$  vertices.*

## 65:16 Optimal Electrical Oblivious Routing on Expanders

Next, we claim that the effective resistance of our demand pairs is the same up to a constant in  $G_{\Delta,n}$  and  $G_{\Delta,n}^\Phi$ .

▷ **Claim 5.7.** For each edge  $(u, v) \in E(G_{\Delta,n})$ , the effective resistance of the pair  $(u, v)$  in the graph  $G_{\Delta,n}^\Phi$  is  $\Theta(1)$ .

*Proof.* To show this result, we give an explicit mapping of the electrical flow routing  $\chi_{(u,v)}$  in  $G_{\Delta,n}$  to routing the flow in  $G_{\Delta,n}^\Phi$  whose energy is at most constant. Let  $\mathbf{f}$  be this electrical flow routing on  $G_{\Delta,n}$ , then we map the flow on each edge  $e'$  in  $G_{\Delta,n}$  uniformly through the  $1/\Phi$  disjoint paths between the endpoints of  $e'$  in  $G_{\Delta,n}^\Phi$ . Since each path now routes only a  $\Phi$ -fraction of the original flow on the edge  $e'$ , we have that the energy used to route through each edge on the disjoint paths replacing  $e'_{middle}$  is  $(\mathbf{f}(e')\Phi)^2 = \mathbf{f}(e')^2\Phi^2$ . We thus have that the energy incurred by routing through the  $1/\Phi$  disjoint paths each consisting of  $1/\Phi$  edges is  $1/\Phi^2 \cdot \mathbf{f}(e')^2\Phi^2 = \mathbf{f}(e')^2$ . Thus, the effective resistance of  $(u, v)$  in  $G_{\Delta,n}^\Phi$  is at most the resistance in  $G_{\Delta,n}$  which implies it is in  $O(1)$ .

A lower bound of  $\Omega(1)$  is observed by inverting this mapping to collect the amount of flow pushed through the disjoint paths replacing edge  $e'$  together and adding it to  $e'$  in  $G_{\Delta,n}$ . The proof is straightforward and therefore omitted. ◁

Before we can carry out the proof of our lower bound, it remains to show for our lower bound gadget which is the graph  $G = G_{\Delta,n} \cup G_{\Delta,n}^\Phi$  that it is a  $\Theta(\Phi)$ -expander.

▷ **Claim 5.8.**  $G$  is  $\Theta(\Phi)$ -expander with  $\Theta(n/\Phi^2)$  edges.

*Proof.* The number of edges is straightforward from our construction of  $G$ . To see that  $G$  is  $O(\Phi)$ -expander, observe that we can take the internal vertices of any path in  $G_{\Delta,n}^\Phi$  replacing an edge in  $G_{\Delta,n}$  which has volume  $\Omega(1/\Phi)$  but only two edges leaving (the once to the endpoints of the replaced edges). To observe that it is an  $\Omega(\Phi)$ -expander, it suffices to show that each cut in  $S$  is maximized by assigning all internal vertices of each such path to one side of the cut. It is then not hard to show from  $G_{\Delta,n}$  being a  $\Omega(1)$ -expander that the claim follows. ◁

Let us now give the proof of the main result. We take the graph under consideration to be  $G = G_{\Delta,n} \cup G_{\Delta,n}^\Phi$ . We take as demand, the vector  $\chi^\Phi = \frac{1}{\Phi} \cdot \chi = \frac{1}{\Phi} \cdot \{\chi_{(u,v)}\}_{e=(u,v) \in E(G_{\Delta,n})}$ . Let  $\mathbf{A}_\mathcal{E}$  denote the electrical flow routing on this graph  $G$ . Let us look at each edge  $e = (u, v) \in E(G_{\Delta,n})$ . From Fact 5.3, Claim 5.7 and Fact 5.4, we have that the flow  $\mathbf{f}_e = \mathbf{A}_\mathcal{E} \cdot \frac{1}{\Phi} \chi_{(u,v)}$  restricted to the edges in  $E(G_{\Delta,n})$  routes in total at least  $\log n / \varphi$  units of flow along all of these edges. By linearity of  $\mathbf{A}_\mathcal{E}$  and the fact that flows do not cancel, we have that when routing  $\chi^\Phi$ , an average edge in  $E(G_{\Delta,n})$  carries  $\Omega(\log n / \Phi)$  units of flow. Thus, the  $\ell_p$ -norm of this flow is at least  $\sqrt[p]{n \cdot (\log n / \Phi)^p} = n^{1/p} \cdot \log n \cdot \Phi$ . But observe that we can route the flow with demand with congestion 1 in  $G_{\Delta,n}^\Phi$  by routing for each demand  $\chi_{(u,v)}$  exactly 1 unit of flow through each of the disjoint paths corresponding to the edge  $e' = (u, v)$  in  $G_{\Delta,n}^\Phi$ . The  $\ell_p$ -norm of this flow is  $\Theta((n/\Phi^2)^{1/p}) = \Theta(n^{1/p} \Phi^{-2/p})$  (using Claim 5.8). We thus have that  $\rho_p(\mathbf{A}_\mathcal{E}) = \Omega(\log n \cdot \Phi^{p/(p-2)})$ .

To obtain the result for  $p < 2$ , we use that for  $q$  given by  $1/p + 1/q = 1$ , we have  $\|\mathbf{X}\|_{p \rightarrow p} = \|\mathbf{X}^\top\|_{q \rightarrow q}$ , and for the electrical routing,  $\mathbf{A}_\mathcal{E} \mathbf{B} = (\mathbf{A}_\mathcal{E} \mathbf{B})^\top$  since  $\mathbf{L}^+$  is symmetric.

We note that in the construction above the number of vertices in the final graph  $G$  might be much larger than  $n$ . By considering all possible parameters for  $\Phi$  in  $[1/n, 1]$  (i.e. all such numbers were  $1/\Phi$  is integer), we obtain a family of  $n'$ -vertex graphs with conductances in  $[1, 1/\sqrt[3]{n'}]$ , as claimed. Since every  $\Phi$ -expander is also a  $\Phi'$ -expander for every  $\Phi' \leq \Phi$ , we do further not need to restrict the domain of  $\Phi'$  further than in range. We point out that by considering parameters  $\Phi$  in our construction that are even smaller than  $1/n$ , one can get up to an arbitrarily small polynomial factor close to conductances as small as  $1/\sqrt{n'}$ .



## References

- 1 N Alon and V. D Milman.  $\Lambda_1$ , Isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38(1):73–88, February 1985. doi:10.1016/0095-8956(85)90092-9.
- 2 Noga Alon, Shirshendu Ganguly, and Nikhil Srivastava. High-girth near-ramanujan graphs with localized eigenvectors. *Israel Journal of Mathematics*, 246(1):1–20, 2021.
- 3 Yossi Azar, Edith Cohen, Amos Fiat, Haim Kaplan, and Harald Racke. Optimal oblivious routing in polynomial time. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 383–388, 2003.
- 4 Yair Bartal and Stefano Leonardi. On-line routing in all-optical networks. In Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Automata, Languages and Programming*, Lecture Notes in Computer Science, pages 516–526, Berlin, Heidelberg, 1997. Springer. doi:10.1007/3-540-63165-8\_207.
- 5 Jeff Cheeger. A lower bound for the smallest eigenvalue of the Laplacian, *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*, 1970.
- 6 Li Chen, Rasmus Kyng, Yang P. Liu, Simon Meierhans, and Maximilian Probst Gutenberg. Almost-linear time algorithms for incremental graphs: Cycle detection, sccs, s-t shortest path, and minimum-cost flow. *CoRR*, abs/2311.18295, 2023. To appear at STOC’24. doi:10.48550/arXiv.2311.18295.
- 7 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum Flow and Minimum-Cost Flow in Almost-Linear Time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623, October 2022. doi:10.1109/FOCS54457.2022.00064.
- 8 Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. Solving sdd linear systems in nearly  $m \log^{1/2} n$  time. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 343–352, 2014.
- 9 Matthias Englert and Harald Räcke. Oblivious routing for the  $L_p$ -norm. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 32–40. IEEE, 2009.
- 10 Mohsen Ghaffari, Bernhard Haeupler, and Goran Zuzic. Hop-constrained oblivious routing. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1208–1220, New York, NY, USA, June 2021. Association for Computing Machinery. doi:10.1145/3406325.3451098.
- 11 Gramoz Goranci, Monika Henzinger, Harald Räcke, Sushant Sachdeva, and A. R. Sricharan. Electrical flows for polylogarithmic competitive oblivious routing, 2023. arXiv:2303.02491.
- 12 Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. The expander hierarchy and its applications to dynamic graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2212–2228. SIAM, 2021.
- 13 Bernhard Haeupler, Harald Räcke, and Mohsen Ghaffari. Hop-constrained expander decompositions, oblivious routing, and distributed universal optimality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 1325–1338, New York, NY, USA, June 2022. Association for Computing Machinery. doi:10.1145/3519935.3520026.
- 14 Arun Jambulapati and Aaron Sidford. Ultrasparse ultrasparsifiers and faster laplacian system solvers. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 540–559. SIAM, 2021.
- 15 Jonathan Kelner and Petar Maymounkov. Electric routing and concurrent flow cutting. *Theoretical computer science*, 412(32):4123–4135, 2011.
- 16 Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations, 2013. arXiv:1304.2338.
- 17 Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving sdd systems in nearly-linear time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 911–920, 2013.

- 18 Rasmus Kyng, Richard Peng, Sushant Sachdeva, and Di Wang. Flows in almost linear time via adaptive preconditioning. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 902–913, 2019.
- 19 Rasmus Kyng and Maximilian Probst. Advanced graph algorithms and optimization, 2021.
- 20 Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 573–582. IEEE, 2016.
- 21 Gregory Lawler and Hariharan Narayanan. Mixing times and  $\ell_p$  bounds for oblivious routing. In *2009 Proceedings of the Sixth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pages 66–74. SIAM, 2009.
- 22 Huan Li and Aaron Schild. Spectral subspace sparsification. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 385–396. IEEE, 2018.
- 23 Jason Li. Faster parallel algorithm for approximate shortest path. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 308–321, 2020.
- 24 Aleksander Madry. Fast Approximation Algorithms for Cut-Based Problems in Undirected Graphs. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 245–254, October 2010. doi:10.1109/FOCS.2010.30.
- 25 B.M. Maggs, F. Meyer auf der Heide, B. Vocking, and M. Westermann. Exploiting locality for data management in systems of limited bandwidth. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 284–293, October 1997. doi:10.1109/SFCS.1997.646117.
- 26 Richard Peng. Approximate undirected maximum flows in  $o(m \text{ polylog}(n))$  time. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1862–1867. SIAM, 2016.
- 27 H. Racke. Minimizing congestion in general networks. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 43–52, November 2002. doi:10.1109/SFCS.2002.1181881.
- 28 Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, pages 255–264, 2008.
- 29 Harald Räcke, Chintan Shah, and Hanjo Täubig. Computing Cut-Based Hierarchical Decompositions in Almost Linear Time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 227–238. Society for Industrial and Applied Mathematics, January 2014. doi:10.1137/1.9781611973402.17.
- 30 Václav Rozhoň, Christoph Grunau, Bernhard Haeupler, Goran Zuzic, and Jason Li. Undirected  $(1 + \epsilon)$ -shortest paths via minor-aggregates: near-optimal deterministic parallel and distributed algorithms. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 478–487, 2022.
- 31 Aaron Schild. An almost-linear time algorithm for uniform random spanning tree generation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 214–227, 2018.
- 32 Aaron Schild, Satish Rao, and Nikhil Srivastava. Localization of electrical flows. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 1577–1584, USA, January 2018. Society for Industrial and Applied Mathematics.
- 33 Jonah Sherman. Nearly maximum flows in nearly linear time. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 263–269. IEEE, 2013.
- 34 Jonah Sherman. Generalized preconditioning and undirected minimum-cost flow. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 772–780. SIAM, 2017.
- 35 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 81–90, New York, NY, USA, June 2004. Association for Computing Machinery. doi:10.1145/1007352.1007372.

- 36 Elias M. Stein and Guido Weiss. *The Fourier Transform*, pages 1–36. Princeton University Press, 1971. URL: <http://www.jstor.org/stable/j.ctt1bpm9w6.4>.
- 37 L. G. Valiant and G. J. Brebner. Universal schemes for parallel communication. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*, STOC '81, pages 263–277, New York, NY, USA, 1981. Association for Computing Machinery. doi:10.1145/800076.802479.
- 38 Leslie G. Valiant. A scheme for fast parallel communication. *SIAM journal on computing*, 11(2):350–361, 1982.
- 39 Goran Zuzic, Gramoz Goranci, Mingquan Ye, Bernhard Haeupler, and Xiaorui Sun. Universally-optimal distributed shortest paths and transshipment via graph-based  $\ell_1$ -oblivious routing. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 2549–2579. SIAM, 2022. doi:10.1137/1.9781611977073.100.