

MIT Open Access Articles

Electric routing and concurrent flow cutting

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Kelner, Jonathan, and Petar Maymounkov. "Electric Routing and Concurrent Flow Cutting." Algorithms and Computation. Ed. Yingfei Dong, Ding-Zhu Du, & Oscar Ibarra. (Lecture Notes in Computer Science ; Vol. 5878). Berlin, Heidelberg: Springer Berlin Heidelberg, 2009. 792–801. Web.

As Published: http://dx.doi.org/10.1007/978-3-642-10631-6_80

Publisher: Springer-Verlag

Persistent URL: <http://hdl.handle.net/1721.1/71604>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Electric routing and concurrent flow cutting

Jonathan Kelner *

Petar Maymounkov *

September 15, 2009

Abstract: We investigate an oblivious routing scheme, amenable to distributed computation and resilient to graph changes, based on electrical flow. Our main technical contribution is a new rounding method which we use to obtain a bound on the $\ell_1 \rightarrow \ell_1$ operator norm of the inverse graph Laplacian. We show how this norm reflects both latency and congestion of electric routing.

1 Introduction

Overview We address a vision of the Internet where every participant exchanges messages with their direct friends and no one else. Yet such an Internet should be able to support reliable and efficient routing to remote locations identified by unchanging names in the presence of an ever changing graph of connectivity.

Modestly to this end, this paper investigates the properties of routing along the electric flow in a graph (*electric routing* for short) for intended use in such distributed systems, whose topology changes over time. We focus on the class of expanding graphs which, we believe, gives a good trade-off between applicability and precision in modeling a large class of real-world networks. We address distributed representation and computation. As a measure of performance, we show that electric routing, being an oblivious routing scheme, achieves minimal maximum edge congestion (as compared to a demand-dependent optimal scheme). Furthermore, we show that electric routing continues to work (on average) in the presence of large edge failures occurring after the routing scheme has been computed, which attests to its applicability in changing environments. We now proceed to a formal definition of oblivious routing and statement of our results.

*Research partially supported by NSF grant CCF-0843915.

ACM Classification: F.2.2, G.2.2

AMS Classification: 68Q25, 68R10

Key words and phrases: oblivious routing, spectral graph theory, Laplacian operator

Oblivious routing The object of interest is a graph $G = (V, E)$ (with $V = [n]$ and $|E| = m$) undirected, positively edge-weighted by $w_{u,v} \geq 0$, and not necessarily simple. The intention is that higher $w_{u,v}$ signifies stronger connectivity between u and v ; in particular, $w_{u,v} = 0$ indicates the absence of edge (u, v) . For analysis purposes, we fix an arbitrary orientation “ \rightarrow ” on the edges (u, v) of G , i.e. if (u, v) is an edge then exactly one of $u \rightarrow v$ or $v \rightarrow u$ holds. Two important operators are associated to every G . The *discrete gradient* operator $B \in \mathbf{R}^{E \times V}$, sending functions on V to functions on the *undirected* edge set E , is defined as $\chi_{(u,v)}^* B := \chi_u - \chi_v$ if $u \rightarrow v$, and $\chi_{(u,v)}^* B := \chi_v - \chi_u$ otherwise, where χ_y is the Kronecker delta function with mass on y . For $e \in E$, we use the shorthand $B_e := (\chi_e B)^*$. The *discrete divergence* operator is defined as B^* .

A (*single-commodity*) *demand* of amount $\alpha > 0$ between $s \in V$ and $t \in V$ is defined as the vector $d = \alpha(\chi_s - \chi_t) \in \mathbf{R}^V$. A (*single-commodity*) *flow* on G is defined as a vector $f \in \mathbf{R}^E$, so that $f_{(u,v)}$ equals the flow value from u towards v if $u \rightarrow v$, and the negative of this value otherwise. We also use the notation $f_{u \rightarrow v} := f_{(u,v)}$ if $u \rightarrow v$, and $f_{u \rightarrow v} := -f_{(u,v)}$ otherwise. We say that flow f *routes* demand d if $B^* f = d$. This is a linear algebraic way of encoding the fact that f is an (s, t) -flow of amount α . A *multi-commodity demand*, also called a *demand set*, is a matrix whose columns constitute the individual demands’ vectors. It is given as the direct product $\oplus_\tau d_\tau$ of its columns. A *multi-commodity flow* is represented as a matrix $\oplus_\tau f_\tau$, given as a direct product of its columns, the single-commodity flows. For clarity, we write $f_{\tau,e}$ for $(f_\tau)_e$. The flow $\oplus_\tau f_\tau$ *routes* the demand set $\oplus_\tau d_\tau$ if $B^* f_\tau = d_\tau$, for all τ , or in matrix notation $B^*(\oplus_\tau f_\tau) = \oplus_\tau d_\tau$. The *congestion* $\|\cdot\|_G$ of a multi-commodity flow measures the load of the most-loaded edge, relative to its capacity. It is given by

$$\|\oplus_\tau f_\tau\|_G := \max_e \sum_\tau |f_{\tau,e}/w_e| = \|(\oplus_\tau f_\tau)^* W^{-1}\|_{1 \rightarrow 1}, \text{ where } \|A\|_{1 \rightarrow 1} := \sup_{x \neq 0} \frac{\|Ax\|_1}{\|x\|_1}. \quad (1.1)$$

An *oblivious routing scheme* is a (not necessarily linear) function $R : \mathbf{R}^V \rightarrow \mathbf{R}^E$ which has the property that $R(d)$ routes d when d is a valid single-commodity demand (according to our definition). We extend R to a function over demand sets by defining $R(\oplus_\tau d_\tau) := \oplus_\tau R(d_\tau)$. This says that each demand in a set is routed independently of the others by its corresponding R -flow. We measure the “goodness” of an oblivious routing scheme by the maximum traffic that it incurs on an edge (relative to its capacity) compared to that of the optimal (demand-dependent) routing. This is captured by the *competitive ratio* η_R of the routing scheme R , defined

$$\eta_R := \sup_{\oplus_\tau d_\tau} \sup_{\substack{\oplus_\tau f_\tau \\ B^*(\oplus_\tau f_\tau) = \oplus_\tau d_\tau}} \frac{\|R(\oplus_\tau d_\tau)\|_G}{\|\oplus_\tau f_\tau\|_G}. \quad (1.2)$$

Let \mathcal{E} denote the (yet undefined) function corresponding to electric routing. Our main theorem states:

THEOREM 1.1 — *For every undirected graph G with unit capacity edges, maximum degree d_{\max} and vertex expansion $\alpha := \min_{S \subseteq V} \frac{|E(S, S^c)|}{\min\{|S|, |S^c|\}}$, one has $\eta_{\mathcal{E}} \leq \left(4 \ln \frac{n}{2}\right) \cdot \left(\alpha \ln \frac{2d_{\max}}{2d_{\max} - \alpha}\right)^{-1}$. This is tight up to a factor of $O(\ln \ln n)$.*

The competitive ratio in Theorem 1.1 is best achievable for any oblivious routing scheme up to a factor of $O(\ln \ln n)$ due to a lower bound for expanders, i.e. the case $\alpha = O(1)$, given in [7]. Theorem 1.1 can be extended to other definitions of graph expansion, weighted and unbounded-degree graphs. We omit these extensions for brevity. We also give an unconditional, albeit much worse, bound on $\eta_{\mathcal{E}}$:

THEOREM 1.2 — *For every unweighted graph on m edges, electrical routing has $\eta_{\mathcal{E}} \leq O(m^{1/2})$. Furthermore, there are families of graphs with corresponding demand sets for which $\eta_{\mathcal{E}} = \Omega(m^{1/2})$.*

Electric routing Let $W = \text{diag}(\dots, w_e, \dots) \in \mathbf{R}^{E \times E}$ be the edge weights matrix. We appeal to a known connection between graph Laplacians and electric current [5, 13]. Graph edges are viewed as wires of resistance w_e^{-1} and vertices are viewed as connection points. If $\varphi \in \mathbf{R}^V$ is a vector of vertex potentials then, by Ohm’s law, the *electric flow* (over the edge set) is given by $f = WB\varphi$ and the corresponding demand is $B^*f = L\varphi$ where the (*un-normalized*) Laplacian L is defined as $L = B^*WB$. Central to the present work will be the vertex potentials that induce a desired (s, t) -flow, given by $\varphi^{[s,t]} = L^\dagger(\chi_s - \chi_t)$, where L^\dagger is the pseudo-inverse of L . Thus, the electric flow corresponding to the demand pair (s, t) is $WB\varphi^{[s,t]} = WBL^\dagger(\chi_s - \chi_t)$. We define the *electric routing* operator as

$$\mathcal{E}(d) = WBL^\dagger d \quad (1.3)$$

The vector $\mathcal{E}(\chi_s - \chi_t) \in \mathbf{R}^E$ encodes a unit flow from s to t supported on G , where the flow along an edge (u, v) is given by $\llbracket st, uv \rrbracket := \mathcal{E}(\chi_s - \chi_t)_{u \rightarrow v} = (\varphi_u^{[s,t]} - \varphi_v^{[s,t]})w_{u,v}$.¹ (Our convention is that current flows towards lower potential.) When routing an indivisible message (an IP packet e.g.), we can view the unit flow $\mathcal{E}(\chi_s - \chi_t)$ as a distribution over (s, t) -paths defined recursively as follows: Start at s . At any vertex u , *forward* the message along an edge with positive flow, with probability proportional to the edge flow value. Stop when t is reached. This rule defines the *electric walk* from s to t . It is immediate that the flow value over an edge (u, v) equals the probability that the electric walk traverses that edge.

Let “ \sim ” denote the vertex adjacency relation of G . In order to make a (divisible or indivisible) forwarding decision, a vertex u must be able to compute $\llbracket st, uv \rrbracket$ for all neighbors $v \sim u$ and all pairs $(s, t) \in \binom{V}{2}$. We address this next.

Representation In order to compute $\llbracket st, uv \rrbracket$ (for all $s, t \in V$ and all $v \sim u$) at u , it suffices that u stores the vector $\varphi^{[w]} := L^\dagger\chi_w$, for all $w \in \{w : w \sim u\} \cup \{u\}$. This is apparent from writing

$$\llbracket st, uv \rrbracket = (\chi_u - \chi_v)L^\dagger(\chi_s - \chi_t) = (\varphi^{[u]} - \varphi^{[v]})^*(\chi_s - \chi_t), \quad (1.4)$$

where we have (crucially) used the fact that L^\dagger is symmetric. The vectors $\varphi^{[w]}$ stored at u comprise the (*routing*) *table* of u , which consists of $\deg(u) \cdot n$ real numbers. Thus the per-vertex table sizes of our scheme grow linearly with the vertex degree – a property we call *fair* representation. It seems that fair representation is key for routing in heterogenous systems consisting of devices with varying capabilities.

¹The bilinear form $\llbracket st, uv \rrbracket = \chi_{s,t}BL^\dagger B^*\chi_{u,v}$ acts like a “representation” of G , hence the custom bracket notation.

Equation (1.4), written as $\llbracket st, uv \rrbracket = (\chi_s - \chi_t)^* (\varphi^{[u]} - \varphi^{[v]})$, shows that in order to compute $\llbracket st, uv \rrbracket$ at u , it suffices to know the *indices* of s and t (in the $\varphi^{[w]}$'s). These indices could be represented by $O(\ln n)$ -bit opaque vertex ID's and could be carried in the message headers. Routing schemes that support opaque vertex addressing are called *name-independent*. Name independence allows for vertex name persistence across time (i.e. changing graph topology) and across multiple co-existing routing schemes.

Computation We use an idealized computational model to facilitate this exposition. The vertices of G are viewed as processors, synchronized by a global step counter. During a time step, pairs of processors can exchange messages of arbitrary (finite) size as long as they are connected by an edge. We describe an algorithm for computing approximations $\tilde{\varphi}^{[v]}$ to all $\varphi^{[v]}$ in $O(\ln n/\lambda)$ steps, where λ is the Fiedler eigenvalue of G (the smallest non-zero eigenvalue of L). If G is an expander, then $\lambda = O(1)$. At every step the algorithm sends messages consisting of $O(n)$ real numbers across every edge and performs $O(\deg(v) \cdot n)$ arithmetic operations on each processor v . Using standard techniques, this algorithm can be converted into a relatively easy-to-implement asynchronous one. (We omit this detail from here.) It is assumed that no graph changes occur during the computation of vertex tables.

A vector $\zeta \in \mathbf{R}^V$ is *distributed* if ζ_v is stored at v , for all v . A matrix $M \in \mathbf{R}^{V \times V}$ is *local* (with respect to G) if $M_{u,v} \neq 0$ implies $u \sim v$ or $u = v$. It is straightforward that if ζ is distributed and M is local, then $M\zeta$ can be computed in a single step, resulting in a new distributed vector. Extending this technique shows that for any polynomial $q(\cdot)$, the vector $q(M)\zeta$ can be computed in $\deg(q)$ steps.

The Power Method gives us a matrix polynomial $q(\cdot)$ of degree $O(\ln n/\lambda)$ such that $q(L)$ is a “good” approximation of L^\dagger . We compute the distributed vectors $\zeta^{[w]} := q(L)\chi_w$, for all w , in parallel. As a result, each vertex u obtains $\tilde{\varphi}^{[u]} = (\zeta_u^{[1]}, \dots, \zeta_u^{[n]})$, which approximates $\varphi^{[u]}$ according to Theorem 1.3 and the symmetry of L . In one last step, every processor u sends $\tilde{\varphi}^{[u]}$ to its neighbors. The approximation error n^{-5} is chosen to suffice (in accordance with Corollary 5.2) as discussed next.

THEOREM 1.3 — *Let λ be the Fiedler (smallest non-zero) eigenvalue of G 's Laplacian L , and let G be of bounded degree d_{\max} . Then $\|\zeta^{[v]} - \varphi^{[v]}\|_2 \leq n^{-5}$, where $\zeta^{[v]} = (2d_{\max})^{-1} \sum_{\omega=0}^k M^\omega \chi_v$ and $M = I - L/2d_{\max}$, as long as $k \geq \Omega(\lambda^{-1} \cdot \ln n)$.*

Robustness and latency In order to get a handle on the analysis of routing in an ever-changing network we use a simplifying assumption: the graph does not change during the computation phase while it can change afterwards, during the routing phase. This assumption is informally justified because the computation phase in expander graphs (which we consider to be the typical case) is relatively fast, it takes $O(\ln n)$ steps. The routing phase, on the other hand, should be as “long” as possible before we have to recompute the scheme. Roughly, a routing scheme can be used until the graph changes so much from its shape when the scheme was computed that both the probability of reaching destinations and the congestion properties of the scheme deteriorate with respect to the new shape of the graph. We quantify the robustness of electric routing against edge removals in the following two theorems:

THEOREM 1.4 — *Let G be an unweighted graph with Fiedler eigenvalue $\lambda = \Theta(1)$ and maximum degree d_{\max} , and let $f^{[s,t]}$ denote the unit electric flow between s and t . For any $0 < p \leq 1$, let $Q_p = \{e \in E : |f_e^{[s,t]}| \geq p\}$ be the set of edges carrying more than p flow. Then, $|Q_p| \leq \min\{2/(\lambda p^2), 2d_{\max}\|L^\dagger\|_{1 \rightarrow 1}/p\}$.*

Note that part one of this theorem, i.e. $|Q_p| \leq 2/(\lambda p^2)$, distinguishes electric routing from simple schemes like shortest-path routing. The next theorem studies how edge removals affect demands when “the entire graph is in use:”

THEOREM 1.5 — *Let graph G be unweighted of bounded degree d_{\max} and vertex expansion α . Let f be a routing of the uniform multi-commodity demand set over V (single unit of demand between every pair of vertices), produced by an η -competitive oblivious routing scheme. Then, for any $0 \leq x \leq 1$, removing a x -fraction of edges from G removes at most $x \cdot (\eta \cdot d_{\max} \cdot \ln n \cdot \alpha^{-1})$ -fraction of flow from f .*

The expected number of edges traversed between source and sink reflects the latency of a routing. We establish (Proven in Appendix D):

THEOREM 1.6 — *The latency of every electric walk on an undirected graph of bounded degree d_{\max} and vertex expansion α is at most $O(\min\{m^{1/2}, d_{\max}\alpha^{-2} \ln n\})$.*

Analysis The main hurdle is Theorem 1.1, which we attack in two steps. First, we show that any linear routing scheme R (i.e. scheme for which the operator $R : \mathbf{R}^V \rightarrow \mathbf{R}^E$ is linear) has a distinct worst-case demand set, known as *uniform demands*, consisting of a unit demand between the endpoints of every edge of G . Combining this with the formulaic expression for electric flow (1.3) gives us an operator-based geometric bound for $\eta_\mathcal{E}$, which in the case of a bounded degree graph is simply $\eta_\mathcal{E} \leq O(\|L^\dagger\|_{1 \rightarrow 1})$ where the operator norm $\|\cdot\|_{1 \rightarrow 1}$ is defined by $\|A\|_{1 \rightarrow 1} := \sup_{x \neq 0} \|Ax\|_1 / \|x\|_1$. This is shown in Theorem 3.1. Second, we give a rounding-type argument that establishes the desired bound on $\|L^\dagger\|_{1 \rightarrow 1}$. This argument relies on a novel technique we dub *concurrent flow cutting* and is our key technical contribution. This is done in Theorem 4.1. This concludes the analysis of the congestion properties of electric flow.

The computational procedure for the vertex potentials $\varphi^{[v]}$'s (above) only affords us approximate versions $\tilde{\varphi}^{[v]}$ with ℓ_2 error guarantees. We need to ensure that, when using these in place of the exact ones, all properties of the exact electric flow are preserved. For this purpose, it is convenient to view the electric flow as a distribution over paths (i.e. the electric walk, defined above) and measure the total variation distance between the walks induced by exact and approximate vertex potentials. This is achieved in Theorem 5.1 and Corollary 5.2. It is then easy to verify that any two multi-commodity flows, whose respective individual flows have sufficiently small variation distance, have essentially identical congestion and robustness properties.

Related work Two bodies of prior literature concern themselves with oblivious routing. One focuses on approximating the shortest-path metric [16, 15, 1, 2], the other focuses on approximating the minimal congestion universally across all possible demand sets [12, 8]. The algorithms in these works are essentially best possible in terms of competitive characteristics, however they are not distributed and do not address (competitive) performance in the presence of churn. It is not obvious how to provide efficient distributed variants for these routing schemes that are additionally resistant to churn. The primary reason for this are the algorithmic primitives used. Common techniques are landmark (a.k.a. beacon) selection [16, 15], hierarchical tree-decomposition or tree-packings [12]. These approaches place disproportionately larger importance on “root” nodes, which makes the resulting schemes vulnerable to individual failures. Furthermore, these algorithms require more than (quasi-)linear time (in the centralized model), which translates to prohibitively slow distributed times.

We are aware of one other work in the theoretical literature by Goyal, et al. [6] that relates to efficient and churn-tolerant distributed routing. Motivated by the proliferation of routing schemes for trees, they show that expanders are well-approximated by the union of $O(1)$ spanning trees. However, they do not provide a routing scheme, since routing over unions of trees is not understood.

Concurrently with this paper, Lawler, et al. [10] study just the congestion of electric flow in isolation from other considerations like computation, representation or tolerance to churn. Their main result is a variant of our graph expansion-based bound on $\|L^\dagger\|_{1 \rightarrow 1}$, given by Theorem 4.1. Our approaches, however, are different. We use a geometric approach, compared to a less direct probabilistic one. Our proof exposes structural information about the electric flow, which makes the fault-tolerance of electric routing against edge removal an easy consequence. This is not the case for the proofs found in [10].

Organization Section 2 covers definitions and preliminaries. Section 3 relates the congestion of electric flow to $\|L^\dagger\|_{1 \rightarrow 1}$. Section 4 obtains a bound on $\|L^\dagger\|_{1 \rightarrow 1}$ by introducing the concurrent flow cutting method. Section 5 relates the electric walk to the electric flow and proves (i) stability against perturbation of vertex potentials, (ii) latency bounds and (iii) robustness theorems. Section 6 contains remarks and open problems.

2 Preliminaries

The Spectral Theory of graphs is comprehensively covered in [3]. The object of interest is a graph $G = (V, E)$ (with $V = [n]$ and $|E| = m$) undirected, positively edge-weighted by $w_{u,v} \geq 0$, and not necessarily simple. Whenever we use unweighted graphs, we have $w_{u,v} = 1$ if $u \sim v$ and $w_{u,v} = 0$ otherwise. The Laplacian is positive semi-definite, $L \geq 0$, and thus can be diagonalized as $L = U\Lambda U^*$, where $U \in \mathbf{R}^{n \times n}$ is unitary and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$. By convention, we write $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. For every G , $\lambda_n \leq 2D$, where D is the maximum degree. When G is connected, $\text{Ker}(L) = \mathbb{1}$, and so $LL^\dagger = L^\dagger L = \pi_{\perp \mathbb{1}}$ where π_W denotes projection onto W and L^\dagger denotes the pseudo-inverse of L . On occasion we use $\lambda_{\min} := \lambda_2$ and $\lambda_{\max} := \lambda_n$. The *vertex expansion* of an unweighted G is defined as

$$\alpha := \min_{S \subseteq V} \frac{|E(S, S^c)|}{\min\{|S|, |S^c|\}}.$$

3 The geometry of congestion

Recall that given a multi-commodity demand, electric routing assigns to each demand the corresponding electric flow in G , which we express (1.3) in operator form $\mathcal{E}(\oplus_\tau d_\tau) := WBL^\dagger(\oplus_\tau d_\tau)$. Electric routing is oblivious, since $\mathcal{E}(\oplus_\tau d_\tau) = \oplus_\tau \mathcal{E}(d_\tau)$ ensures that individual demands are routed independently from each other. The first key step in our analysis, Theorem 3.1, entails bounding $\eta_{\mathcal{E}}$ by the $\|\cdot\|_{1 \rightarrow 1}$ matrix norm of a certain natural graph operator on G . This step hinges on the observation that all linear routing schemes have an easy-to-express worst-case demand set:

THEOREM 3.1 — *For every undirected, weighted graph G , let $\Pi = W^{1/2}BL^\dagger B^*W^{1/2}$, then*

$$\eta_{\mathcal{E}} \leq \|W^{1/2}\Pi W^{-1/2}\|_{1 \rightarrow 1}. \quad (3.1)$$

Proof of Theorem 3.1. It is sufficient to consider demand sets that can be routed in G with unit congestion, since both electric and optimal routing scale linearly with scaling the entire demand set. Let $\oplus_\tau d_\tau$ be any demand set, which can be (optimally) routed in G with unit congestion via the multi-commodity flow $\oplus_\tau f_\tau$. Thus, $d_\tau = \sum_e f_{\tau,e} B_e$, for all τ .

The proof involves two steps:

$$\|\mathcal{E}(\oplus_\tau d_\tau)\|_G \stackrel{(i)}{\leq} \|\mathcal{E}(\oplus_e w_e B_e)\|_G \stackrel{(ii)}{=} \|W^{1/2}\Pi W^{-1/2}\|_{1 \rightarrow 1}$$

Step (i) shows that congestion incurred when routing $\oplus_\tau d_\tau$ is no more than that incurred when routing G 's edges, viewed as demands, through G :

$$\begin{aligned} \|\mathcal{E}(\oplus_\tau d_\tau)\|_G &= \|\oplus_\tau \mathcal{E}(d_\tau)\|_G && \text{(i)} \\ &= \|\oplus_\tau \mathcal{E}(\sum_e f_{\tau,e} B_e)\|_G && \text{use } d_\tau = \sum_e f_{\tau,e} B_e \\ &= \|\oplus_\tau \sum_e \mathcal{E}(f_{\tau,e} B_e)\|_G && \text{use } \mathcal{E}(\sum_j d_j) = \sum_j \mathcal{E}(d_j) \\ &\leq \|\oplus_{\tau,e} \mathcal{E}(f_{\tau,e} B_e)\|_G && \text{use } \|\sum_j f_j\|_G \leq \|\oplus_j f_j\|_G \\ &= \|\oplus_e \mathcal{E}(\sum_\tau |f_{\tau,e}| B_e)\|_G && \text{use } \|\oplus_j \alpha_j f\|_G = \|\sum_j |\alpha_j| f\|_G \\ &\leq \|\oplus_e \mathcal{E}(w_e B_e)\|_G && \text{use } \sum_\tau |f_{\tau,e}| \leq w_e \\ &= \|\mathcal{E}(\oplus_e w_e B_e)\|_G \end{aligned}$$

$$\begin{aligned} \|\mathcal{E}(\oplus_e w_e B_e)\|_G &\stackrel{(1.1)}{=} \|\mathcal{E}(\oplus_e w_e B_e)^* W^{-1}\|_{1 \rightarrow 1} \\ &\stackrel{(1.3)}{=} \|W B L^\dagger B^* W W^{-1}\|_{1 \rightarrow 1} = \|W^{1/2} \Pi W^{-1/2}\|_{1 \rightarrow 1}. \end{aligned} \quad (\text{ii}) \quad \blacksquare$$

Remark 3.2 — Note that the proof of step (i) uses only the linearity of \mathcal{E} and so it holds for any linear routing scheme R , i.e. one has $\|R(\oplus_\tau d_\tau)\|_G \leq \|R(\oplus_e w_e B_e)\|_G$.

Using Theorem 3.1, the unconditional upper bound in Theorem 1.2 is simply a consequence of basic norm inequalities. (See Appendix A for a proof.) Theorem 1.1 provides a much stronger bound on $\eta_\mathcal{E}$ when the underlying graph has high vertex expansion. The lower bound in Theorem 1.1 is due to Hajiaghayi, et al. [7]. They show that every oblivious routing scheme is bound to incur congestion of at least $\Omega(\ln n / \ln \ln n)$ on a certain family of expander graphs. The upper bound in Theorem 1.1 follows from Theorem 3.1, Theorem 4.1 and using that $\|\Pi\|_{1 \rightarrow 1} = O(\|L^\dagger\|_{1 \rightarrow 1})$ for unweighted bounded-degree graphs. Thus in the next section we derive a bound on $\|L^\dagger\|_{1 \rightarrow 1}$ in terms of vertex expansion.

4 L_1 operator inequalities

The main results here are an upper and lower bound on $\|L^\dagger\|_{1 \rightarrow 1}$, which match for bounded-degree expander graphs. In this section, we present vertex expansion versions of these bounds that assume bounded-degree.

THEOREM 4.1 — *Let graph $G = (V, E)$ be unweighted, of bounded degree d_{\max} , and vertex expansion*

$$\alpha = \min_{S \subseteq V} \frac{|E(S, S^C)|}{\min\{|S|, |S^C|\}}, \quad \text{then } \|L^\dagger\|_{1 \rightarrow 1} \leq \left(4 \ln \frac{n}{2}\right) \cdot \left(\alpha \ln \frac{2d_{\max}}{2d_{\max} - \alpha}\right)^{-1}. \quad (4.1)$$

The proof of this theorem (given in the next Section) boils down to a structural decomposition of unit (s, t) -electric flows in a graph (not necessarily an expander). We believe that this decomposition is of independent interest. In the case of bounded-degree expanders, one can informally say that the electric walk corresponding to the electric flow between s and t takes every path with probability exponentially inversely proportional to its length. We complement Theorem 4.1 with a lower bound on $\|L^\dagger\|_{1 \rightarrow 1}$ proven in Appendix B:

THEOREM 4.2 — *Let graph $G = (V, E)$ be unweighted, of bounded degree d_{\max} , with metric diameter D . Then, $\|L^\dagger\|_{1 \rightarrow 1} \geq 2D/d_{\max}$ and, in particular, $\|L^\dagger\|_{1 \rightarrow 1} \geq (2 \ln n) \cdot (d_{\max} \ln d_{\max})^{-1}$ for all bounded-degree, unweighted graphs with vertex expansion $\alpha = O(1)$.*

4.1 Proof of upper bound on $\|L^\dagger\|_{1 \rightarrow 1}$ for expanders

Proof of Theorem 4.1. Reformulation: We start by transforming the problem in a more manageable form,

$$\|L^\dagger\|_{1 \rightarrow 1} := \sup_{y \neq 0} \frac{\|L^\dagger y\|_1}{\|y\|_1} = \max_w \|L^\dagger \chi_w\|_1 \stackrel{(*)}{\leq} \frac{n-1}{n} \max_{s \neq t} \|L^\dagger(\chi_s - \chi_t)\|_1, \quad (4.2)$$

where the latter inequality comes from

$$\|L^\dagger \chi_s\|_1 = \|L^\dagger \pi_{\perp 1} \chi_s\|_1 = \|n^{-1} \sum_{t \neq s} L^\dagger(\chi_s - \chi_t)\|_1 \leq \frac{n-1}{n} \max_t \|L^\dagger(\chi_s - \chi_t)\|_1.$$

Pick any vertices $s \neq t$ and set $\psi = L^\dagger(\chi_s - \chi_t)$. In light of (4.2) our goal is to bound $\|\psi\|_1$. We think of ψ as the vertex potentials corresponding to electric flow with imbalance $\chi_s - \chi_t$. By an easy perturbation argument we can assume that no two vertex potentials coincide.

Index the vertices in $[n]$ by increasing potential as $\psi_1 < \dots < \psi_n$. Further, assume that n is even and choose a median c_0 so that $\psi_1 < \dots < \psi_{n/2} < c_0 < \psi_{n/2+1} < \dots < \psi_n$. (If n is odd, then set c_0 to equal the potential of the middle vertex.)

We aim to upper bound $\|\psi\|_1$, given as $\|\psi\|_1 = \sum_v |\psi_v| = \sum_{v: \psi_v > 0} \psi_v - \sum_{u: \psi_u < 0} \psi_u$. Using that $\sum_w \psi_w = 0$, we get $\|\psi\|_1 = 2 \sum_{v: \psi_v > 0} \psi_v = -2 \sum_{u: \psi_u < 0} \psi_u$.

Assume, without loss of generality, that $0 < c_0$, in which case

$$\|\psi\|_1 = -2 \sum_{u: \psi_u < 0} \psi_u \leq 2 \sum_{i=1}^{n/2} |\psi_i - c_0| =: 2N \quad (4.3)$$

In what follows we aim to upper-bound N .

Flow cutting: Define a collection of cuts (S_i, S_i^C) of the form $S_i = \{v : \psi_v \leq c_i\}$, for integers $i \geq 0$, where S_i will be the smaller side of the cut by construction. Let k_i be the number of edges cut by (S_i, S_i^C) and p_{ij} be the length of the j^{th} edge across the same cut. The cut points c_i , for $i \geq 1$, are defined according to $c_i = c_{i-1} - \Delta_{i-1}$, where $\Delta_{i-1} := 2 \sum_j \frac{p_{i-1,j}}{k_{i-1}}$. The last cut, (S_{r+1}, S_{r+1}^C) , is the first cut in the sequence c_0, c_1, \dots, c_{r+1} with $k_{r+1} = 0$ or, equivalently, $S_{r+1} = \emptyset$.

Bound on number of cuts: Let $n_i = |S_i|$. The isoperimetric inequality for vertex expansion (4.1) applied to (S_i, S_i^C) and the fact that $n_i \leq n/2$, by construction, imply

$$\frac{k_i}{n_i} \geq \alpha. \quad (4.4)$$

Let l_i be the number of edges crossing (S_i, S_i^C) that do not extend across c_{i+1} , i.e. edges that are not adjacent to S_{i+1} . The choice $\Delta_i := 2 \sum_j p_{ij}/k_i$ ensures that $l_i \geq k_i/2$. These edges are supported on at

least l_i/d_{\max} vertices in S_i , and therefore $n_{i+1} \leq n_i - l_i/d_{\max}$. Thus,

$$n_{i+1} \leq n_i - \frac{l_i}{d_{\max}} \leq n_i - \frac{k_i}{2d_{\max}} \stackrel{(4.4)}{\leq} n_i - \frac{\alpha n_i}{2d_{\max}} = n_i \left(1 - \frac{\alpha}{2d_{\max}}\right), \quad (4.5)$$

Combining inequality (4.5) with $n_0 = n/2$, we get

$$n_i \leq \frac{n}{2} \left(1 - \frac{\alpha}{2d_{\max}}\right)^i \quad (4.6)$$

The stopping condition implies $S_r \neq \emptyset$, or $n_r \geq 1$, and together with (4.6) this results in

$$r \leq \log_{1/\theta} \frac{n}{2}, \text{ with } \theta = 1 - \frac{\alpha}{2d_{\max}}. \quad (4.7)$$

Amortization argument: Continuing from (4.3),

$$N = \sum_{i=1}^{n/2} |\psi_i - c_0| \stackrel{(*)}{\leq} \sum_{i=0}^r (n_i - n_{i+1}) \sum_{j=0}^i \Delta_j, \quad (4.8)$$

where (*) follows from the fact that for every vertex $v \in S_i - S_{i+1}$ we can write $|\psi_v - c_0| \leq \sum_{j=0}^i \Delta_j$.

Because $BL^\dagger(\chi_s - \chi_t)$ is a unit flow, we have the crucial (and easy to verify) property that, for all i , $\sum_j p_{ij} = 1$. In other words, the total flow of ‘‘simultaneous’’ edges is 1. So,

$$\Delta_i = 2 \sum_j \frac{p_{ij}}{k_i} = \frac{2}{k_i} \stackrel{(4.4)}{\leq} \frac{2}{\alpha n_i} \quad (4.9)$$

Now we can use this bound on Δ_j in (4.8),

$$\sum_{i=0}^r (n_i - n_{i+1}) \sum_{j=0}^i \Delta_j \stackrel{(*)}{=} \sum_{i=1}^r n_i \Delta_i \leq \frac{2}{\alpha} \sum_{i=0}^r 1 = \frac{2}{\alpha} (r + 1),$$

where to derive (*) we use $n_{r+1} = 0$. Combining the above inequality with (4.7) concludes the proof. ■

5 Electric walk

To every unit flow $f \in \mathbf{R}^E$, not necessarily electric, we associate a random walk $W = W_0, W_1, \dots$ called the *flow walk*, defined as follows. Let $\sigma := B^* f$ and so $\sum_v \sigma_v = 0$. The walk starts at W_0 , with

$$\mathbb{P}\{W_0 = v\} = \frac{2 \cdot \max\{0, \sigma_v\}}{\sum_w |\sigma_w|} = \frac{\sum_w f_{v \rightarrow w} - \sum_w f_{w \rightarrow v}}{\sum_u (\sum_w f_{u \rightarrow w} - \sum_w f_{w \rightarrow u})}$$

If the walk is currently at W_t , the next vertex is chosen according to $\mathbb{P}\{W_{t+1} = v \mid W_t = u\} = \frac{f_{u \rightarrow v}}{\sum_w f_{u \rightarrow w}}$, where

$$f_{u \rightarrow v} = \begin{cases} |f_{(u,v)}|, & (u, v) \in E \text{ and } f \text{ flows from } u \text{ to } v \\ 0, & \text{otherwise.} \end{cases} \quad (5.1)$$

When the underlying flow f is an electric flow, i.e. when $f = \mathcal{E}(B^* f)$, the flow walk deserves the specialized name *electric walk*. We study two aspects of electric walks here: (i) stability against perturbations of the vertex potentials, and (ii) robustness against edge removal.

5.1 Stability

The set of vertex potential vectors $\varphi^{[v]} = L^\dagger \chi_v$, for all $v \in V$, encodes all electric flows, as argued in (1.4). In an algorithmic setting, only approximations $\tilde{\varphi}^{[v]}$ of these vectors are available. We ensure that when these approximations are sufficiently good in an ℓ_2 sense, the path probabilities (and congestion properties) of electric walks are virtually unchanged. The next theorem is proven in Appendix C:

THEOREM 5.1 — *Let $\tilde{\varphi}^{[v]}$ be an approximation of $\varphi^{[v]}$, for all $v \in G$, in the sense that*

$$\|\varphi^{[v]} - \tilde{\varphi}^{[v]}\|_2 \leq \nu, \text{ for all } v \in V, \text{ with } \nu = n^{-A}, \quad (5.2)$$

where $A > 4$ is a constant. Then for every electric walk, defined by vertex potentials $\varphi = \sum_v \alpha_v \varphi^{[v]}$, the corresponding “approximate” walk, defined by vertex potentials $\tilde{\varphi} = \sum_v \alpha_v \tilde{\varphi}^{[v]}$, induces a distribution over paths γ with

$$\sum_\gamma \left| \mathbb{P}_\varphi\{W = \gamma\} - \mathbb{P}_{\tilde{\varphi}}\{W = \gamma\} \right| \leq O(n^{2-\frac{A}{2}}), \quad (5.3)$$

where γ ranges over all paths in G , and $\mathbb{P}_\varphi\{W = \gamma\}$ denotes the probability of γ under φ (respectively for $\mathbb{P}_{\tilde{\varphi}}\{W = \gamma\}$).

As shown in Theorem F.1, the Power Method affords us any sufficiently large exponent A , say $A = 5$, without sacrificing efficiency in terms of distributed computation time. In this case, the following corollary asserts that routing with approximate potentials preserves both the congestion properties of the exact electrical flow as well as the probability of reaching the sink.

COROLLARY 5.2 — *Under the assumptions of Theorem 5.1 and $A = 5$, the electric walk defined by vertex potentials $\tilde{\varphi} = \tilde{\varphi}^{[s]} - \tilde{\varphi}^{[t]}$ reaches t with probability $1 - o_n(1)$. Furthermore, for every edge (u, v) with non-negligible load, i.e. $|\tilde{\varphi}_u - \tilde{\varphi}_v| = \omega(n^{-2})$, we have $|\tilde{\varphi}_u - \tilde{\varphi}_v| \rightarrow_n |\varphi_u - \varphi_v|$, where $\varphi = \varphi^{[s]} - \varphi^{[t]}$.*

5.2 Robustness

We prove Theorem 1.4 here, since its proof interestingly relies on the flow cutting techniques developed in this paper. Theorem 1.5 is proved in Appendix E.

Proof of Theorem 1.4. For the first part, let $\{(u_1, v_1), \dots, (u_k, v_k)\} = Q_p$ and let $p_i = |f_{(u_i, v_i)}^{[s, t]}| = |(\chi_{u_i} - \chi_{v_i})^* L^\dagger(\chi_s - \chi_t)|$. Consider the embedding $\zeta : V \rightarrow \mathbf{R}$, defined by $\zeta(v) = \chi_v^* L^\dagger(\delta_s - \delta_t)$. Assume for convenience that $\zeta(u_i) \leq \zeta(v_i)$ for all i . Let $\zeta_{\min} = \min_v \zeta(v)$ and $\zeta_{\max} = \max_v \zeta(v)$.

Choose c uniformly at random from $[\zeta_{\min}, \zeta_{\max}]$ and let $X_i = p_i \cdot \mathbf{I}\{\zeta(u_i) \leq c \leq \zeta(v_i)\}$, where $\mathbf{I}\{\cdot\}$ is the indicator function. Observe that the random variable $X = \sum_i X_i$ equals the total electric flow of all edges in Q_p cut by c . Since these edges are concurrent (in the electric flow) by construction, we have $X \leq 1$. On the other hand,

$$\mathbf{E}X = \sum_i p_i \cdot \mathbb{P}\{\zeta(u_i) \leq c \leq \zeta(v_i)\} \geq \sum_i p_i \frac{p_i}{\zeta_{\max} - \zeta_{\min}} \geq \sum_i \frac{\lambda p_i^2}{2} \geq k \frac{\lambda p^2}{2}$$

Combining this with $\mathbf{E}X \leq 1$ produces $|Q_p| \leq 2/(\lambda p^2)$.

For the second part, $k p \leq \sum_{e \in Q_p} |f_e^{[s, t]}| \leq \sum_{e \in E} |f_e^{[s, t]}| = \|BL^\dagger(\chi_s - \chi_t)\|_1 \leq 2d_{\max} \cdot \|L^\dagger\|_{1 \rightarrow 1}$. This gives $|Q_p| \leq 2d_{\max} \cdot \|L^\dagger\|_{1 \rightarrow 1}/p$. \blacksquare

6 Conclusions

Our main result in Theorem 1.1 attests to the good congestion properties on graphs of bounded degree and high vertex expansion, i.e. $\alpha = O(1)$. A variation on the proof of this theorem establishes a similar bound on $\eta_{\mathcal{E}}$, however, independent of the degree bound and as a function of the *edge expansion* $\beta = \min_{S \subseteq V} \frac{\text{vol}(E(S, S^c))}{\min\{\text{vol}(S), \text{vol}(S^c)\}}$, where $\text{vol}(S) := \sum_{v \in S} \sum_{u: u \sim v} w_{u, v}$ and $\text{vol}(E(S, S^c)) := \sum_{(u, v): u \in S, v \in S^c} w_{u, v}$.

The bounded degree assumption is also implicit in our computational procedure in that all vertices must know an upper bound on d_{\max} in order to apply M in Theorem 1.3. Using a generous bound, anything $\omega(1)$, on d_{\max} is bad because it slows down the mixing of the power polynomial. To avoid this complication, one must use a symmetrization trick outlined in Appendix G.

We conclude with a couple of open questions. A central concern, widely-studied in social-networks, are Sybil Attacks [4]. These can be modeled as graph-theoretic noise, as defined in [9]. It is interesting to understand how such noise affects electric routing. We suspect that any $O(\ln n)$ -competitive oblivious routing scheme, which outputs its routes in the “next hop” model, must maintain $\Omega(n)$ -size routing tables at every vertex. In the *next hop* model, every vertex v must be able to answer the question “What is the flow of the (s, t) -route in the neighborhood of v ?” in time $O(\text{polylog}(n))$, using its own routing table alone and for every source-sink pair (s, t) .

References

- [1] I. Abraham, C. Gavoille, and D. Malkhi. On space-stretch trade-offs: Lower bounds. In *Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures*, pages 207–216. ACM New York, NY, USA, 2006. 6

- [2] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. 2008. 6
- [3] F.R.K. Chung. *Spectral graph theory*. American Mathematical Society, 1997. 6, 20
- [4] J.R. Douceur. The sybil attack. In *Peer-To-Peer Systems: First International Workshop, Iptps 2002, Cambridge, Ma, USA, March 7-8, 2002, Revised Papers*, page 251. Springer, 2002. 12
- [5] P.G. Doyle and J.L. Snell. Random walks and electric networks. *Arxiv preprint math.PR/0001057*, 2000. 3
- [6] N. Goyal, L. Rademacher, and S. Vempala. Expanders via random spanning trees. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 576–585. Society for Industrial and Applied Mathematics Philadelphia, PA, USA, 2009. 6
- [7] M.T. Hajiaghayi, R.D. Kleinberg, T. Leighton, and H. Räcke. New lower bounds for oblivious routing in undirected graphs. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 918–927. ACM New York, NY, USA, 2006. 3, 8
- [8] P. Harsha, T.P. Hayes, H. Narayanan, H. Räcke, and J. Radhakrishnan. Minimizing average latency in oblivious routing. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 200–207. Society for Industrial and Applied Mathematics Philadelphia, PA, USA, 2008. 6
- [9] S. Kale, Y. Peres, and C. Seshadhri. Noise Tolerance of Expanders and Sublinear Expander Reconstruction. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science-Volume 00*, pages 719–728. IEEE Computer Society Washington, DC, USA, 2008. 12
- [10] G. Lawler and H. Narayanan. Mixing times and l_p bounds for Oblivious routing. In *Workshop on Analytic Algorithmics and Combinatorics,(ANALCO09)*. 6
- [11] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999. 19
- [12] H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 255–264. ACM New York, NY, USA, 2008. 6
- [13] D.A. Spielman. *Graphs and networks, Lecture Notes*. 3
- [14] D.A. Spielman and N. Srivastava. Graph sparsification by effective resistances. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 563–568. ACM New York, NY, USA, 2008. 14

- [15] M. Thorup and U. Zwick. Compact routing schemes. In *Proceedings of the thirteenth annual ACM symposium on Parallel algorithms and architectures*, pages 1–10. ACM New York, NY, USA, 2001. 6
- [16] M. Thorup and U. Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005. 6

A Proof of Theorem 1.2

Proof of Theorem 1.2. The upper bound follows from:

$$\eta \stackrel{(3.1)}{\leq} \|\Pi\|_{1 \rightarrow 1} \stackrel{(A.2)}{\leq} m^{1/2} \cdot \|\Pi\|_{2 \rightarrow 2} \stackrel{(A.3)}{=} m^{1/2} \quad (\text{A.1})$$

The second step is justified as follows

$$\|\Pi\|_{1 \rightarrow 1} = \max_e \|\Pi \chi_e\|_1 \leq m^{1/2} \cdot \max_e \|\Pi \chi_e\|_2 \leq m^{1/2} \cdot \|\Pi\|_{2 \rightarrow 2}. \quad (\text{A.2})$$

The third step is the assertion

$$\|\Pi\|_{2 \rightarrow 2} \leq 1, \quad (\text{A.3})$$

which follows from the (easy) fact that Π is a projection, shown by Spielman, et al. in Lemma A.1.

The lower bound is achieved by a graph obtained by gluing the endpoints of \sqrt{n} copies of a path of length \sqrt{n} and a single edge. Routing a flow of value \sqrt{n} between these endpoints incurs congestion $\sqrt{n}/2$. ■

LEMMA A.1 (*Proven in [14]*) — Π is a projection; $\text{Im}(\Pi) = \text{Im}(W^{1/2}B)$; The eigenvalues of Π are 1 with multiplicity $n - 1$ and 0 with multiplicity $m - n + 1$; and $\Pi_{e,e} = \|\Pi \chi_e\|^2$.

B Proof of Theorem 4.2

Proof of Theorem 4.2. Let $s \neq t$ be a pair of vertices in G at distance D . We consider the flow $f = BL^\dagger(\chi_s - \chi_t)$. Set $\psi = L^\dagger(\chi_s - \chi_t)$, and note that we can use $\|f\|_1$ as a lower bound on $\|L^\dagger\|_{1 \rightarrow 1}$,

$$\|f\|_1 = \sum_{(u,v)} |\psi_u - \psi_v| \leq d_{\max} \sum_v |\psi_v| \leq d_{\max} \|L^\dagger(\chi_s - \chi_t)\|_1 \leq \frac{d_{\max}}{2} \|L^\dagger\|_{1 \rightarrow 1}.$$

Now, let $\{\pi_i\}_i$ be a path decomposition of f and let $l(\pi_i)$ and $f(\pi_i)$ denote the length and value, respectively, of π_i . Then,

$$\|f\|_1 = \sum_{(u,v)} |\psi_u - \psi_v| = \sum_i l(\pi_i) f(\pi_i) \geq D \sum_i f(\pi_i) = D. \quad \blacksquare$$

C Proof of Theorem 5.1

Proof of Theorem 5.1. Notation: Note that in this proof we use the notation of (5.1). So, for a potential vector ψ , we have $\psi_{u \rightarrow v} := \psi_u - \psi_v$ if (u, v) is an edge and $\psi_u \geq \psi_v$, and $\psi_{u \rightarrow v} := 0$ otherwise. So, for example, the potential difference on (u, v) can be written as $\psi_{u \rightarrow v} + \psi_{v \rightarrow u}$. On the other hand, we use the single letter edge notation φ_e to denote the signed (according to B) potential difference on e , so $\varphi_e := \varphi_u - \varphi_v$ if $B_e = \delta_u - \delta_v$. Let D be the maximum degree.

Edge approximation: Fix any unit electric flow, defined by potentials $\varphi := \sum_v \alpha_v \varphi^{[v]}$, and write its approximation as $\tilde{\varphi} := \sum_v \alpha_v \tilde{\varphi}^{[v]}$. All unit flows can be so expressed under the restriction that $\sum_v \alpha_v = 0$ and $\sum_v |\alpha_v| = 2$. The approximation condition (5.2) combined with Lemma C.1 gives us, for every edge $e = (u, v)$,

$$\begin{aligned} |\varphi_e - \tilde{\varphi}_e| &= \left| \sum_v \alpha_v (\varphi_e^{[v]} - \tilde{\varphi}_e^{[v]}) \right| \\ &\leq \sum_v |\alpha_v| \cdot |\varphi_e^{[v]} - \tilde{\varphi}_e^{[v]}| \\ &\leq \sum_v |\alpha_v| \cdot 2v && \text{apply Lemma C.1} \\ &= 4v \end{aligned}$$

We call this the *additive edge approximation condition*

$$\varphi_e - 4v \leq \tilde{\varphi}_e \leq \varphi_e + 4v \quad (\text{C.1})$$

Now, consider a fixed path γ along the electric flow defined by φ , traversing vertices w_0, w_1, \dots, w_k . Let $\mathbb{P}_\varphi\{W = \gamma\}$ and $\mathbb{P}_{\tilde{\varphi}}\{W = \gamma\}$ denote the probability of this path under the potentials φ and $\tilde{\varphi}$, respectively. In most of what follows, we build machinery to relate one to the other.

Path probabilities: For a general unit flow (not necessarily an (s, t) -flow), defined by vertex potentials ψ , $\mathbb{P}_\psi\{W = \gamma\}$ equals

$$\mathbb{P}_\psi\{W_0 = w_0\} \left(\prod_{i=0}^{k-1} \mathbb{P}_\psi\{W_{i+1} = w_{i+1} \mid W_i = w_i\} \right) \mathbb{P}_\psi\{W_\infty = w_k \mid w_k\}, \quad (\text{C.2})$$

where next we explain each factor in turn.

The first, $\mathbb{P}_\psi\{W_0 = w_0\}$, is the probability that the walk starts from w_0 , and is expressed as

$$\mathbb{P}_\psi\{W_0 = w_0\} = \max \left(0, \sum_u \psi_{w_0 \rightarrow u} - \sum_u \psi_{u \rightarrow w_0} \right). \quad (\text{C.3})$$

The second and trickiest, $\mathbb{P}_\psi\{W_{i+1} = w_{i+1} \mid W_i = w_i\}$, is the probability that having reached w_i the

walk traverses the edge leading to w_{i+1} , and $\sum_u \psi_{w_i \rightarrow u} \geq \sum_u \psi_{u \rightarrow w_i}$, we write

$$\mathbb{P}_\psi\{W_{i+1} = w_{i+1} \mid W_i = w_i\} = \frac{\psi_{w_i \rightarrow w_{i+1}}}{\max\left(\sum_u \psi_{u \rightarrow w_i}, \sum_u \psi_{w_i \rightarrow u}\right)}. \quad (\text{C.4})$$

To grasp the meaning of the denominator, note that the quantity $|\sum_u \psi_{u \rightarrow w_i} - \sum_u \psi_{w_i \rightarrow u}|$ is the magnitude of the in or out flow (depending on the case) at w_i .

The third, $\mathbb{P}_\psi\{W_\infty = w_k \mid w_k\}$, is the probability that the walk ends (or exits) at w_k conditioned on having reached w_k , and

$$\mathbb{P}_\psi\{W_\infty = w_k \mid w_k\} = \max\left(0, \sum_u \psi_{u \rightarrow w_k} - \sum_u \psi_{w_k \rightarrow u}\right). \quad (\text{C.5})$$

Next, we are going to find multiplicative bounds for all three factors by focusing on ‘‘dominant’’ paths, and discarding ones with overall negligible probability.

Dominant paths: It is straightforward to verify (from first principles) that the probability that an edge (u, v) occurs in the electric walk equals $|\varphi_e| = \varphi_{u \rightarrow v} + \varphi_{v \rightarrow u}$. We call an edge *short* if $|\varphi_e| \leq \epsilon$, where the exact asymptotic of $\epsilon > 0$ is determined later, but for the moment $\nu \ll e \ll 1$. We restrict our attention to *dominant* paths γ that traverse no short edges, and have $\mathbb{P}_\varphi\{W_0 = w_0\} \geq \epsilon$ and $\mathbb{P}_\varphi\{W_\infty = w_k \mid w_k\} \geq \epsilon$.

Indeed, by a union bound, the probability that the electric walk traverses a non-dominant path is at most $2n\epsilon + n^2\epsilon$. This will be negligible and such paths will be of no interest. In summary,

$$\mathbb{P}_\varphi\{W \text{ dominant}\} \geq 1 - 2n\epsilon - n^2\epsilon \quad (\text{C.6})$$

We now condition on the event that γ is dominant.

The no short edge condition gives $\epsilon \leq |\varphi_e| \leq 1$, and using (C.1) we derive the stronger *multiplicative edge approximation condition*

$$\frac{1}{\sigma} \leq \frac{\tilde{\varphi}_e}{\varphi_e} \leq \sigma, \text{ where } \sigma = 1 + \frac{8D\nu}{\epsilon}, \quad (\text{C.7})$$

which holds as long as $\epsilon \geq 4\nu$, as guaranteed by the asymptotics of ϵ . Also note that the latter condition ensures that φ_e and $\tilde{\varphi}_e$ have the same sign. An extra factor of $2D$ is included in σ with foresight.

For the first factor (C.3), we have

$$\begin{aligned} \mathbb{P}_{\tilde{\varphi}}\{W_0 = w_0\} &= \sum_u \tilde{\varphi}_{w_0 \rightarrow u} - \sum_u \tilde{\varphi}_{u \rightarrow w_0} \\ &\geq \sum_u \varphi_{w_0 \rightarrow u} - \sum_u \varphi_{u \rightarrow w_0} - 4D\nu && \text{use (C.1)} \\ &\geq \mathbb{P}_\varphi\{W_0 = w_0\} \left(1 - \frac{4D\nu}{\epsilon}\right) && \text{use } \mathbb{P}_\varphi\{W_0 = w_0\} \geq \epsilon \\ &\geq \frac{1}{\sigma} \mathbb{P}_\varphi\{W_0 = w_0\} && \text{use } \epsilon \leq 1/2. \end{aligned} \quad (\text{C.8})$$

For the second factor (C.4), assume $\sum_u \psi_{u \rightarrow w_i} \geq \sum_u \psi_{w_i \rightarrow u}$. An identical argument holds in the other case. Abbreviate

$$\mathbb{P}_{\tilde{\varphi}}\{w_{i+1} | w_i\} := \mathbb{P}_{\tilde{\varphi}}\{W_{i+1} = w_{i+1} | W_i = w_i\}.$$

Path dominance implies $\sum_u \varphi_{u \rightarrow w_i} \geq \epsilon$, and so

$$\begin{aligned} \mathbb{P}_{\tilde{\varphi}}\{w_{i+1} | w_i\} &= \frac{\tilde{\varphi}_{w_i \rightarrow w_{i+1}}}{\sum_u \tilde{\varphi}_{u \rightarrow w_i}} && \text{(C.9)} \\ &\geq \frac{\sigma^{-1} \cdot \varphi_{w_i \rightarrow w_{i+1}}}{\sum_u \varphi_{u \rightarrow w_i} + 4D\nu} && \text{use (C.7) and (C.1)} \\ &\geq \sigma^{-2} \frac{\varphi_{w_i \rightarrow w_{i+1}}}{\sum_u \varphi_{u \rightarrow w_i}} && \text{use } \epsilon \leq 1/2 \text{ and } \sum_u \varphi_{u \rightarrow w_i} \geq \epsilon \\ &= \frac{1}{\sigma^2} \mathbb{P}_{\varphi}\{w_{i+1} | w_i\}. \end{aligned}$$

For the third factor (C.5), similarly to the first, we have

$$\begin{aligned} \mathbb{P}_{\tilde{\varphi}}\{W_{\infty} = w_k | w_k\} &= && \text{(C.10)} \\ &= \sum_u \tilde{\varphi}_{u \rightarrow w_k} - \sum_u \tilde{\varphi}_{w_k \rightarrow u} \\ &\geq \sum_u \varphi_{u \rightarrow w_k} - \sum_u \varphi_{w_k \rightarrow u} - 4D\nu && \text{use (C.1)} \\ &\geq \mathbb{P}_{\varphi}\{W_{\infty} = w_k | w_k\} \left(1 - \frac{4D\nu}{\epsilon}\right) && \text{use } \mathbb{P}_{\varphi}\{W_{\infty} = w_k | w_k\} \geq \epsilon \\ &\geq \frac{1}{\sigma} \mathbb{P}_{\varphi}\{W_{\infty} = w_k | w_k\} && \text{use } \epsilon \leq 1/2. \end{aligned}$$

Dominant path bound: We now obtain a relation between $\mathbb{P}_{\varphi}\{W = \gamma\}$ and $\mathbb{P}_{\tilde{\varphi}}\{W = \gamma\}$ by combining the bounds (C.8), (C.9) and (C.10) with (C.2):

$$\begin{aligned} \frac{\mathbb{P}_{\tilde{\varphi}}\{W = \gamma\}}{\mathbb{P}_{\varphi}\{W = \gamma\}} &\geq \frac{1}{\sigma^{2n+2}} && \text{apply bounds, and path length } \leq n \text{ (C.11)} \\ &\geq \left(1 - \frac{8D\nu}{\epsilon}\right)^{2n+2} && \text{use } \sigma^{-1} \geq 1 - 8D\nu/\epsilon \\ &\geq \exp\left(-\frac{16D\nu}{\epsilon}\right)^{2n+2} && \text{use } 1 - x \geq e^{-2x} \\ &=: \theta \end{aligned}$$

Statistical difference: Abbreviate $p(\gamma) := \mathbb{P}_{\varphi}\{W = \gamma\}$ and $q(\gamma) := \mathbb{P}_{\tilde{\varphi}}\{W = \gamma\}$. Below, γ iterates through all paths, ζ iterates through dominant paths and ξ iterates through non-dominant paths. We

bound the statistical difference (5.3), using (C.11) which says $q(\zeta) \geq \theta \cdot p(\zeta)$,

$$\begin{aligned}
 \sum_{\gamma} |p(\gamma) - q(\gamma)| &= \\
 &= \sum_{\zeta} |p(\zeta) - q(\zeta)| + \sum_{\xi} |p(\xi) - q(\xi)| \\
 &\leq \sum_{\zeta} |(1 - \theta)p(\zeta) - (q(\zeta) - \theta p(\zeta))| + \sum_{\xi} q(\xi) && \text{use } \theta < 1 \\
 &\leq \sum_{\zeta} |q(\zeta) - \theta p(\zeta)| + \sum_{\xi} q(\xi) \\
 &= 1 - \sum_{\zeta} p(\zeta) \\
 &= (1 - \theta) + \theta \sum_{\zeta} p(\zeta) \tag{C.12}
 \end{aligned}$$

In this final step, we pin-point the asymptotics of ϵ that simultaneously minimize the two terms of (C.12). In the following, we parameterize $\epsilon = n^{-B}$ and use (C.6),

$$\begin{aligned}
 (1 - \theta) + \theta \sum_{\zeta} p(\zeta) &= \\
 &= 1 - \exp\left(-\frac{16Dv}{\epsilon}\right)^{2n+2} + \exp\left(-\frac{16Dv}{\epsilon}\right)^{2n+2} (2n\epsilon + n^2\epsilon) \\
 &= 1 - \exp O(-Dn^{B-A+1}) + n^{2-B} \cdot \exp O(-Dn^{B-A+1}) \\
 &= O(Dn^{B-A+1}) + n^{2-B} \cdot \exp O(-Dn^{B-A+1}), \quad \text{use } 1 - e^{-x} \leq x \\
 &= O(n^{B-A+2}) + O(n^{2-B}) \quad \text{use } D \leq n \\
 &= O(n^{2-\frac{A}{2}}), \quad \text{set } B = A/2. \quad \blacksquare
 \end{aligned}$$

LEMMA C.1 — *If $x, y \in \ell_2$ and $\|x - y\|_2 \leq v$, then for all $i \neq j$,*

$$(x_i - x_j) - 2v \leq y_i - y_j \leq (x_i - x_j) + 2v.$$

Proof. We have $(x_i - y_i)^2 \leq \|x - y\|_2^2 \leq v^2$, implying $|x_i - y_i| \leq v$. Similarly for j . Combining the two proves the lemma. \blacksquare

D Proof of Theorem 1.6

Proof of Theorem 1.6. Let $X_e^{[s,t]}$ be the indicator that edge e participates in the electric walk between s and t . Then the latency can be expressed as

$$\begin{aligned} \max_{s \neq t} \sum_e \mathbf{E} X_e^{[s,t]} &= \max_{s \neq t} \sum_{(u,v)} |(\delta_u - \delta_v) L^\dagger (\delta_s - \delta_t)| \\ &= \max_{s \neq t} \|BL^\dagger (\delta_s - \delta_t)\|_1 = \|BL^\dagger B^*\|_{1 \rightarrow 1} = \|\Pi\|_{1 \rightarrow 1}. \end{aligned}$$

The latter is bounded by Theorem 4.1 and $\|\Pi\|_{1 \rightarrow 1} \leq m^{1/2}$, as in (A.1) e.g. ■

Remark D.1 — For expanders, this theorem is not trivial. In fact, there exist path realizations of the electric walk which can traverse up to $O(n)$ edges. Theorem 1.6 asserts that this happens with small probability. On the other hand, in a bounded-degree expander, even if s and t are adjacent the walk will still take a $O(\log n)$ -length path with constant probability.

E Proof of Theorem 1.5

Proof of Theorem 1.5. Let f_{opt} be a max-flow routing of the uniform demands and let θ be the fraction of the demand set that is routed by f_{opt} . The Multi-commodity Min-cut Max-flow Gap Theorem (Theorem 2, in [11]) asserts

$$O(\ln n) \cdot \theta \geq \min_{S \subset V} \frac{|E(S, S^c)|}{|S| \cdot |S^c|} \geq \frac{1}{n} \cdot \min_{S \subset V} \frac{|E(S, S^c)|}{\min\{|S|, |S^c|\}} = \frac{\alpha}{n}$$

Thus the total demand flown by f_{opt} is no less than $\theta \binom{n}{2} \geq \Omega(\alpha n / \ln n)$. Normalize f (by scaling) so it routes the same demands as f_{opt} . If k edges are removed, then at most ηk flow is removed from f , which is at most a fraction $\eta k \cdot O(\ln n / \alpha n)$ of the total flow. Substitute $x = k/m$ and use $m \leq d_{\max} n$ to complete the proof. ■

F Proof of Theorem 1.3

Theorem 1.3 is implied by the following theorem by specializing $\epsilon = O(n^{-5})$:

THEOREM F.1 — *Let G be a graph, whose Laplacian L has smallest eigenvalue λ and whose maximum degree is D . Then, for every y with $\|y\|_2 = 1$ the vector $x = L^\dagger y$ can be approximated using*

$$\tilde{x} = \frac{1}{2D} \sum_{i=0}^d \left(I - \frac{L}{2D}\right)^i y,$$

so that for every $\epsilon > 0$,

$$\|x - \tilde{x}\|_2 \leq \epsilon, \text{ as long as } d \geq \Omega(1) \cdot \ln \frac{1}{\lambda \epsilon D} \cdot \left(\ln \frac{1}{1 - \lambda} \right)^{-1}.$$

Proof of Theorem F.1. We normalize L via $N = L/\tau$ (and so $L^{-1} = N^{-1}/\tau$), where $\tau = 2D$. Since $\tau = 2D \geq \lambda_{\max}(L)$, the eigenvalues of N are in $[0, 1]$. In this case, the Moore-Penrose inverse of N is given by $N^\dagger = \sum_{i=0}^{\infty} (I - N)^i$. Set $N_0^\dagger = \sum_{i=0}^d (I - N)^i$ and $N_1^\dagger = N^\dagger - N_0^\dagger$. Our aim is to minimize d so that

$$\|x - \tilde{x}\|_2 = \left\| \frac{N_0^\dagger + N_1^\dagger}{\tau} y - \frac{N_0^\dagger}{\tau} y \right\|_2 = \left\| \frac{N_1^\dagger}{\tau} y \right\|_2 \leq \left\| \frac{N_1^\dagger}{\tau} \right\|_{2 \rightarrow 2} \leq \epsilon,$$

where $\|A\|_{2 \rightarrow 2} := \sup_{x \neq 0} \|Ax\|_2 / \|x\|_2$ denotes the matrix spectral norm. Set $\kappa := \tau / \lambda_{\min}$, so that κ^{-1} is the smallest eigenvalue of N ,

$$\begin{aligned} \|N_1^\dagger\|_{2 \rightarrow 2} &= \left\| \sum_{i=d+1}^{\infty} (I - N)^i \right\|_{2 \rightarrow 2} \leq \sum_{i=d+1}^{\infty} \|(I - N)^i\|_{2 \rightarrow 2} \\ &\leq \sum_{i=d+1}^{\infty} (1 - \kappa^{-1})^i = (1 - \kappa^{-1})^{d+1} \kappa \end{aligned} \quad (\text{F.1})$$

Setting (F.1) less than $\tau\epsilon$ gives

$$d \geq \frac{\ln \kappa / (\tau\epsilon)}{\ln \kappa / (\kappa - 1)}. \quad \blacksquare$$

G Symmetrized algorithm

In this section we discuss how to modify the computational procedure, given in the Section 1, in order to apply it to graphs of unbounded degree. The described algorithm for computing $\varphi^{[w]} = L^\dagger \chi_w$ relies on the approximation of L^\dagger via the Taylor series $\frac{1}{1-x} = \sum_{i=0}^{\infty} (1-x)^i$. The series converges only when $\|x\|_2 < 1$, which is ensured by setting $x = \frac{L}{2d_{\max}}$, and using that $\|L\|_{2 \rightarrow 2} < 2d_{\max}$. Thus we arrive at $2d_{\max} \cdot L^\dagger = \sum_{i=0}^{\infty} (I - \frac{L}{2d_{\max}})^i$. This approach continues to work if we replace d_{\max} with any upper bound $h_{\max} \geq d_{\max}$, obtaining $L^\dagger = \frac{1}{2h_{\max}} \sum_{i=0}^{\infty} (I - M)^i$ where $M = \frac{L}{2h_{\max}}$, however this is done at the expense of slower convergence of the series. Since in a distributed setting all vertices must agree on what M is, a worst-case upper bound $h_{\max} = n$ must be used, which results in a prohibitively slow convergence even for expander graphs.

Instead, we pursue a different route. Let $\mathcal{L} = D^{-1/2} L D^{-1/2}$ be the *normalized Laplacian* of G , where $D \in \mathbf{R}^{n \times n}$ is diagonal with $D_{v,v} = \deg(v)$. One always has $\|\mathcal{L}\|_{2 \rightarrow 2} \leq 2$ (Lemma 1.7 in [3]) while at the same time $\lambda_{\min}(\mathcal{L}) \geq \max\{\frac{\beta^2}{2}, \frac{\alpha^2}{4d_{\max} + 2d_{\max}\alpha}\}$ (Theorems 2.2 and 2.6 in [3]), where α and β are the vertex- and edge-expansion of G , respectively. Set $M = \mathcal{L}/3$, so that $\|M\|_{2 \rightarrow 2} < 1$. Recall that the aim of our

distributed procedure is to compute $\varphi_u^{[w]}$ at u (for all w). We achieve this using the following:

$$\varphi_u^{[w]} = \chi_u^* L^\dagger \chi_w = \chi_u^* D^{-1/2} \frac{M^\dagger}{3} D^{-1/2} \chi_w = \frac{\chi_u^*}{\sqrt{\deg(u)}} \frac{\sum_{i=0}^{\infty} (I - M)^i}{3} \frac{\chi_w}{\sqrt{\deg(w)}}$$

The key facts about the series in the left-hand side are that (i) it converges quickly when G is an expander and (ii) all vertices can compute M locally, in particular, without requiring any global knowledge like e.g. an upper bound on d_{\max} .