Mixing times and ℓ_p bounds for Oblivious routing

Gregory Lawler
Department of Mathematics
University of Chicago *

Hariharan Narayanan
Department of Computer Science
University of Chicago †

Abstract

We study the task of uniformly minimizing all the ℓ_p norms of the vector of edge loads in an undirected graph while obliviously routing a multicommodity flow. Let G be an undirected graph having m edges and n vertices. Let the performance index π of an oblivious routing algorithm $\mathcal{A}$ (on G) be the supremum of its competitive ratios over all ℓ_p norms, where the adversarial adaptive routing scheme may vary both with norm and the set of demands. We give a expression in closed form for π(Harmonic) for a certain oblivious algorithm Harmonic that uses the "electrical flow". We show that $\pi(\text{Harmonic}) = O(T_{mix})$ where $T_{mix}(G)$ is the mixing time of the canonical random walk on G and by $O(\sqrt{m})$. These results lead to $O(\log n)$ upper bounds on π(Harmonic) for expanders. We independently show that on $N \times M$ discrete tori where $N \leq M$, $\pi(\text{Harmonic}) = O(N)$.

Lastly, we can handle a larger class of norms than ℓ_p, namely those norms that are invariant with respect to all permutations of the canonical basis vectors and reflections about any of them. Two novel aspects of our proofs are the use of interpolation theorems that relate different operator norms, and connections between discrete harmonic functions and random walks.

1 Introduction

Over the past three decades, there has been significant interest in the design and analysis of routing schemes in networks of various kinds. A network is typically modeled as a directed or undirected graph $G = (V, E)$, where E is a set of m edges representing links and V is a set of n vertices representing locations or servers. Each link is associated with a cost which is a function of the load that it carries. There is a set of demands, which has the form

$$\{(i,\, j,\, d_{ij}) \big| (i,\, j) \in V \times V,\, d_{ij} \geq 0\}.$$

A routing scheme that routes a commodity from its source to its target independent of the demands at other source-target pairs is termed oblivious. The competitive ratio of an oblivious routing scheme Obl is the maximum taken over all demands, of cost that Obl incurs divided by the cost that the optimal adaptive scheme Opt incurs. Work in this area was initiated by Valiant and Brebner [18] who developed an oblivious routing protocol for parallel routing in the hypercube that routes any permutation in time that is only a logarithmic factor away from optimal. For the cost-measure of congestion (the maximum load of a network link) in a virtual circuit routing model, Räcke [11] proved the existence of an oblivious routing scheme with polylogarithmic competitive ratio for any undirected network. This result was subsequently made constructive by Harrelson, Hildrum and Rao [5] and improved to a competitive ratio of $O(\log^2 n \log\log n)$.

Oblivious routing has largely been studied in the context of minimizing the maximum congestion. A series of papers [11, 5, 12] has culminated recently in the development of an oblivious algorithm due to Räecke [12] whose competitive ratio with respect to congestion is $O(\log n)$. The algorithm Harmonic that is studied in this paper was introduced in [7] where it was shown to have a competitive ratio of $O(\sqrt{\log n})$ with respect to the ℓ_2 norm when demands route to a single common target. We study the task of uniformly minimizing all the ℓ_p norms of the vector of edge loads in an undirected graph while routing a multicommodity flow oblivious of the set of demands. As a matter of fact our results hold under any norm that transforms $\mathbb{R}^n$ into a Banach space symmetric and "unconditional" with respect to the canonical basis. These terms have been defined in section 3.

2 Our results

Let $G = (V, E)$ denote an undirected graph with a set V of n vertices and a set E of m edges. For any oblivious algorithm $\mathcal{A}$, let the competitive ratio of $\mathcal{A}$ in the norm $\|\cdot\|_p$ be denoted $\kappa_p(\mathcal{A})$. Let the performance index $\pi(\mathcal{A})$ of $\mathcal{A}$ be defined to be their

*Research supported by National Science Foundation grant DMS-0734151.

†Research partially supported by William Eckhardt Graduate fellowship

supremum as $\|\cdot\|_p$ ranges over the set of all ℓ_p norms.

$$\pi(\mathcal{A}) := \sup_p \kappa_p(\mathcal{A}).$$

Let Harmonic be the oblivious algorithm (formally defined in the following section,) which routes a flow from s to t in the (unique) way that minimizes the ℓ_2 norm of edge loads of that flow, assuming all other demands to be 0. The competitive ratio of this algorithm with respect to the ℓ_2 norm was shown in [7] to be $O(\sqrt{\log n})$ over demands having single common target. We show that Harmonic has an index π(Harmonic) that is equal to its competitive ratio in the ℓ_1 norm, which is in turn bounded above by $\min(\sqrt{m}, O(T_{mix}))$ where T_{mix} is the mixing time of the canonical random walk, We obtain $O(\log n)$ upper bounds on π(Harmonic) for expanders and two dimension tori. The constant in $O(\cdot)$ may depend on the family. Almost matching $\Omega(\frac{\log n}{\log\log n})$ lower bounds for expanders [6] and matching $\Omega(\log n)$ lower bounds for 2−dimensional discrete tori [2] are known for the competitive ratio of an oblivious algorithm with respect to congestion or ℓ_∞ norm. In particular, for cost functions that are convex combinations of bounded powers of the various ℓ_p norms, such as $\sum_e g(\text{load}(e))$ where g is a polynomial with non-negative coefficients, Harmonic has on these graphs a polylogarithmic competitive ratio. We show that there exist graphs for which no algorithm that is adaptive with respect to the demands but not p can simultaneously have a cost that is less than $\Omega(\sqrt{m})$ times the ℓ_p norm of the optimal adaptive algorithm that is permitted to vary with p, even if p can only take the values 1 and ∞. Lastly, we can handle a larger class of norms than ℓ_p, namely those norms that are invariant with respect to all permutations of the canonical basis vectors and reflections about any of them.

Theorem 2.1. *For any graph G, with n vertices, and m edges, on which the canonical random walk has a mixing time T_{mix}, π(Harmonic) $\leq \min(\sqrt{m}, O(T_{mix}))$.*

Hajiaghayi et al have shown in [6] that if G belongs to a family of expanders and $\mathcal{A}$ is any oblivious routing algorithm, the competitive ratio $\kappa_\infty(\mathcal{A})$ with respect to congestion is bounded from below by $\Omega(\frac{\log n}{\log\log n})$. Therefore Theorem 2.1 is tight up to an $O(\log\log n)$ factor for expanders.

Theorem 2.2. *If G is a two dimensional $N \times M$ torus where $N \leq M$,*

$$\pi(\text{Harmonic})[G] = O(\log N).$$

For 2-dimensional $N \times N$ grids, $\kappa_\infty(\mathcal{A})$ was shown to be bounded below by $\Omega(\log N)$ by Bartal and Leonardi. A minor variation of the same argument gives an $\Omega(\log N)$ lower bound for 2−dimensional $N \times N$ tori as well. Therefore Theorem 2.2 is tight up to a universal constant for square tori.

Theorem 2.3. *For every m, there exists a graph G with m edges, such that for any oblivious algorithm $\mathcal{A}$, $\pi(\mathcal{A}) \geq \frac{\lfloor\sqrt{m-1}\rfloor}{2}$ on G.*

3 Definitions and Preliminaries

A network will be an undirected graph $G = (V, E)$, where V denotes a set of n vertices (or nodes) $\{1, \dots, n\}$ and E a set of m edges. If a traffic vector $t = (t_1, \dots, t_m)$ is transported across edges $e_1, \dots, e_m$, we shall consider costs that are ℓ_p norms $\|t\|_p$. In our setting, the network is undirected and links are allowed to carry traffic in both directions simultaneously. For book keeping, it will be convenient to give each edge an orientation. For an edge e of the form $\{v, w\}$, we will write $e = (v, w)$ when we want to emphasize that the edge is oriented from v to w. The traffic on edge e will be a real number. If this number is positive, it will represent traffic along e from v to w; if it is negative, it will represent traffic along e from w to v. Let $\text{In}(v)$ be the edges of G that are oriented into v and $\text{Out}(v)$ be the edges of G that are oriented away from v. A potential ϕ on G is a function from V to $\mathbb{R}$. The gradient $\nabla\phi$ of a potential ϕ is a function from E to $\mathbb{R}$, whose value on an oriented edge $e := (u, v)$ is

$$\nabla\phi(e) := \phi(u) - \phi(v).$$

A flow f on G is a function from E to $\mathbb{R}$. The divergence $\text{div}\, f$ of a flow f is a function from V to $\mathbb{R}$ whose value on a vertex v is given by

$$(\text{div}\, f)(v) := \sum_{e \in \text{Out}(v)} f(e) - \sum_{e \in \text{In}(v)} f(e)$$

We shall denote by Δ the Laplacian operator that maps the space of real-valued functions on V to itself as follows.

$$\Delta\phi := -\text{div}\,(\nabla\phi).$$

We call such $f = \langle f_{ij} : i, j \in V(G)\rangle$ a multi-flow. We say that a multi-flow f meets the demand $\langle d_{ij} : i, j \in V\rangle$, if for all $i, j \in V$,

$$\text{div}\, f_{ij} = d_{ij}\delta_i - d_{ij}\delta_j,$$

where $\delta_u(\cdot)$ is the Kronecker Delta function that takes a value 1 on u and 0 on all other vertices. If this is the case, we say "f routes D" and write $f \searrow D$. For a fixed i, j, we shall use $\|f_{ij}\|_1$ to denote $\sum_e |f_{ij}(e)|$. The traffic on the edge e under f is given by

$$t_f(e) = \sum_{i,j} |f_{ij}(e)|.$$

We shall call the vector $t_f := (t_f(e_1), \ldots, t_f(e_m))$ the network traffic or network load, where $(e_1, \ldots, e_m)$ is a list of the edges of G. If for every edge e, the total traffic on e under f is greater or equal to the total traffic on e under f', we shall say that $f' \lhd f$. i. e.

$$(\forall e) t_f(e) \geq t_{f'}(e) \Rightarrow f' \lhd f.$$

Definition 1. *An oblivious algorithm ($\mathcal{A}$), is a multi-flow $\{a_{ij}\}$ indexed by pairs of vertices i, j where each a_{ij} is a flow satisfying*

$$\operatorname{div} a_{ij} = \delta_i - \delta_j.$$

Given a demand D, $\mathcal{A}$ routes D using $D \cdot a := \langle d_{ij} a_{ij} : i, j \in V(G) \rangle$.

Definition 2. *For any oblivious algorithm $\mathcal{A}$, for every $p \in [1, \infty]$, we define its competitive ratio under the ℓ_p norm $\| \cdot \|_p$*

$$\kappa_p(\mathcal{A}) := \sup_D \sup_{f \searrow D} \frac{\|t_{D \cdot a}\|_p}{\|t_f\|_p}.$$

Let the performance index $\pi(\mathcal{A})$ of $\mathcal{A}$ be defined to be their supremum as $\| \cdot \|_p$ ranges over all possible ℓ_p norms,

$$\pi(\mathcal{A}) := \sup_{p \in [1, \infty]} \kappa_p(\mathcal{A}).$$

All results in this paper hold without modification if the above definition of performance index, is altered to be the supremum over all norms that satisfy the symmetry and unconditionality conditions in Definition 6.

Definition 3. *We define* Harmonic *to be the oblivious algorithm corresponding to the multi-flow $h = \langle h_{ij} : i, j \in V(G) \rangle$, where h_{ij} is the unique flow such that*

1. *$\operatorname{div} h_{ij} = \delta_i - \delta_j$, and*
2. *There exists a potential ϕ_{ij} such that $\nabla \phi_{ij} = h_{ij}$.*

These conditions uniquely determine $\{h_{ij}\}$ and determine the potential ϕ_{ij} up to an additive constant. The potential can be described in terms of random walk on the graph. Suppose $W_0, W_1, \ldots$ denotes simple random walk on the graph, and let $\tilde{\pi}(v) = \deg(v)/[2\#(E)]$ denotes its stationary distribution.

Definition 4. (Hitting time) *If $S \subseteq V$, let*

$$H_S = \min\{j \geq 1 : W_j \in S\}, \quad \overline{H}_S = \min\{j \geq 0 : W_j \in S\}.$$

If $S = \{i, j\}$, we write $H_{ij}, \overline{H}_{ij}$.

Note that $H_S, \overline{H}_S$ agree if $W_0 \notin S$. The potential ϕ_{ij} with boundary condition $\phi_{ij}(j) = 0$ is given by

$$\phi_{ij}(v) = b_{ij} \, \mathbb{P}^v\{W(\overline{H}_{ij}) = i\},$$

where we write $\mathbb{P}^v$ to denote probabilities assuming $W_0 = v$ and the constant b_{ij} is given by

$$b_{ij}^{-1} = \mathbb{P}^j\{W(H_{ij}) = i\}.$$

Definition 5. (Mixing time) *Let $W_0, W_1, \ldots$ be simple random walk on G. Let*

$$\rho_v^{(t)}(u) = \mathbb{P}^v\{W_t = u\}.$$

The mixing time as a function of ϵ is

$$T_{mix}(\epsilon) := \sup_{v \in V} \inf \left\{ t : \|\tilde{\pi} - \rho_v^{(t)}\|_1 \leq 2\epsilon \right\}.$$

Definition 6. *A Banach space X with a basis $\{e_1, \ldots, e_m\}$ is said to be symmetric and unconditional with respect to the basis if the following two conditions hold for any $x_1, \ldots, x_m \in \mathbb{R}$.*

- *S. For any permutation π, $\|\sum_{i=1}^m x_i e_i\|_X = \|\sum_{i=1}^m x_i e_{\pi(i)}\|_X$.*
- *U. For any $\epsilon_1, \ldots, \epsilon_m \in \{-1, 1\}$, $\|\sum x_i e_i\|_X = \|\sum \epsilon_i x_i e_i\|_X$.*

3.1 Interpolation Theorems All ℓ_p norms satisfy the above conditions. Given a linear operator $A : \mathbb{R}^m \mapsto \mathbb{R}^m$, we shall define its $\ell_p \to \ell_p$ norm

$$\|A\|_{p \to p} = \sup_{\|x\|_p = 1} \frac{\|Ax\|_p}{\|x\|_p}.$$

More generally if $\mathbb{R}^m$ is endowed with a norm transforming it into a Banach space X, we shall denote its operator norm by $\|A\|_{X \to X}$. We will need the following special cases of the theorems of Riesz-Thorin [13, 16] and Mityagin [9].

Theorem 3.1. (Riesz-Thorin) *For any $1 \leq p \leq r \leq q \leq \infty$,*

$$\|A\|_{r \to r} \leq \max(\|A\|_{p \to p}, \|A\|_{q \to q}).$$

The following theorem is due to B. Mityagin.

Theorem 3.2. (Mityagin) *Let $\mathbb{R}^m$ be endowed with a norm transforming it into a Banach space X that is symmetric and unconditional with respect to the standard basis. Then,*

$$\|A\|_{X \to X} \leq \max(\|A\|_{1 \to 1}, \|A\|_{\infty \to \infty}).$$

3.2 Some facts about harmonic functions and flows

$h_{ij} = \nabla\phi_{ij}$, where ϕ_{ij} is up to addition by a constant, the unique solution of the linear equation

$$\Delta\phi_{ij} = -\delta_i + \delta_j.$$

Therefore for every u, v and w, the following linear relation is true.

FACT 3.1. $h_{uv} + h_{vw} = h_{uw}$.

For $e = (u, v)$, let $h_e := h_{uv}$ and $\phi_e := \phi_{uv}$. More generally, we have the following.

LEMMA 3.1. *Let g_{ij} be a flow such that* $\operatorname{div} g_{ij} = \delta_i - \delta_j$. *Then,*

$$\sum_e g_{ij}(e)h_e = h_{ij}.$$

Proof. By linearity,

$$\begin{aligned}\operatorname{div}\sum_e g_{ij}(e)h_e &= \sum_e g_{ij}(e)\operatorname{div} h_e \\ &= \sum_{e=(u,v)\in E} g_{ij}(e)(\delta_u - \delta_v),\end{aligned}$$

which is

$$\sum_{v\in V}\left(\sum_{e\in \mathrm{Out}(v)} g_{ij}(e) - \sum_{e\in \mathrm{In}(v)} g_{ij}(e)\right)\delta_v = \delta_i - \delta_j.$$

Secondly,

$$\nabla(\sum_e g_{ij}(e)\phi_e) = \sum_e g_{ij}(e)\nabla\phi_e = \sum_e g_{ij}(e)h_e.$$

According to the definition, h_{ij} is the unique flow that satisfies the above properties, so we are done.

The following is a result from network theory [17].

THEOREM 3.3. (RECIPROCITY THEOREM) *The flows comprising* HARMONIC *have the following symmetry property. For each $i, j \in V$, let ϕ_{ij} be a potential such that $\nabla\phi_{ij} = h_{ij}$. Then, for any $u, v \in V$,*

$$\phi_{ij}(u) - \phi_{ij}(v) = \phi_{uv}(i) - \phi_{uv}(j). \tag{3.1}$$

4 Using Interpolation to derive uniform bounds

Proof. [Proof of Theorem 2.1] This Theorem follows from Proposition 4.3, Proposition 4.1 and Proposition 4.2.

PROPOSITION 4.1. *For any graph G,*

$$\begin{aligned}\pi(\text{HARMONIC}) &= \kappa_1(\text{HARMONIC}) \\ &= \max_{e\in G}\|h_e\|_1,\end{aligned}$$

where the maximum is taken over all edges of G.

Proof. [Proof of Proposition 4.1] Given a demand D, let D_p be constructed as follows. Let $\text{OPT}_p(G, D)$ be an optimal multi-flow routing D with respect to ℓ_p, and $\text{opt}_p(G, D)$ be the corresponding ℓ_p norm. For an edge $e = (u, w)$, let $(D_p)_{uw}$ be the total amount of traffic from u to v in $\text{OPT}_p(G, D)$ that routes D. For any pair of vertices (u, w) that are not adjacent, let $(D_p)_{uw}$ be defined to be 0. Let $\|D\|_p$ be defined in the natural way to be

$$\|D\|_p := \left(\sum_{ij} d_{ij}^p\right)^{\frac{1}{p}}.$$

LEMMA 4.1. $\text{opt}_p(G, D) = \text{opt}_p(G, D_p) = \|D_p\|_p$.

Proof. [Proof of Lemma 4.1] Any multi-flow f that routes D_p can be converted to a multi-flow that routes D having the same total cost, since D_p was constructed from a multi-flow that routes D. Therefore $\text{opt}_p(G, D) \le \text{opt}_p(G, D_p)$. By the definition of D_p, there exists an optimal solution to D that can be used to route D_p. This establishes that $\|D_p\|_p = \text{opt}_p(G, D) \ge \text{opt}_p(G, D_p)$, and proves the lemma.

Let $\bar{\mathcal{D}}_p$ represent the set of all demands of the form D_p arising from some demand D by the above conversion procedure. By Lemma 4.1,

$$\forall_{D_p\in\bar{\mathcal{D}}_p} \sup_{f\searrow D}\frac{1}{\|t_f\|_p} = \frac{1}{\|D_p\|_p}.$$

Using $D \cdot h$ to denote the multi-flow $\langle d_{ij}h_{ij} : i, j \in V(G)\rangle$, it is sufficient to prove that for all $p \in [1, \infty]$.

$$\sup_{D_p\in\bar{\mathcal{D}}_p}\frac{\|t_{D_p\cdot h}\|_p}{\|D_p\|_p} \le \sup_{D_1\in\bar{\mathcal{D}}_1}\frac{\|t_{D_1\cdot h}\|_1}{\|D_1\|_1}. \tag{4.2}$$

Let $e_1, \ldots, e_m$ be some enumeration of the edges of G and let R be the $m \times m$ matrix whose ij^{th} entry, where $e_i = (u, v)$ is given by $r_{ij} = |h_{uv}(e_j)|$. For $D_p \in \bar{\mathcal{D}}_p$, the "traffic vector" $t_{D_p\cdot h}$ can be obtained by applying the linear transformation R to d_p, for each $e_i = (u, v)$, $(d_p)_i = (D_p)_{uv}$ and $d_p = ((d_p)_1, \ldots, (d_p)_m)$.

$$\begin{pmatrix} t_{D_p\cdot h}(e_1) \\ t_{D_p\cdot h}(e_2) \\ \vdots \\ t_{D_p\cdot h}(e_m)\end{pmatrix} = \begin{pmatrix} r_{11} & \ldots & r_{1m} \\ r_{21} & \ldots & r_{2m} \\ \vdots & \ddots & \vdots \\ r_{m1} & \ldots & r_{mm}\end{pmatrix}\begin{pmatrix}(d_p)_1 \\ (d_p)_2 \\ \vdots \\ (d_p)_m\end{pmatrix}.$$

Given a linear operator $A : \mathbb{R}^m \mapsto \mathbb{R}^m$, we define its $\ell_p \to \ell_p$ norm

$$\|A\|_{p\to p} = \sup_{\|x\|_p=1}\frac{\|Ax\|_p}{\|x\|_p}.$$

With this notation, for $p \in (1, \infty]$,

$$\sup_{D_p \in \bar{\mathcal{D}}_p} \frac{\|t_{D_p \cdot h}\|_p}{\|D_p\|_p} \quad \leq \quad \|R\|_{p \to p} \tag{4.3}$$

while for $p = 1$, because $\bar{\mathcal{D}}_1$ is equal to $\{\mathbb{R}^+\}^m$, we can make the stronger assertion that

$$\sup_{D_1 \in \bar{\mathcal{D}}_1} \frac{\|t_{D_1 \cdot h}\|_1}{\|D_1\|_1} \quad = \quad \|R\|_{1 \to 1}. \tag{4.4}$$

Lemma 4.2. $\|R\|_{1\to 1} = \|R\|_{\infty \to \infty}$.

Proof. [Proof of Lemma 4.2] Recall that R is an $m \times m$ matrix whose ij^{th} entry is $|h_{e_i}(e_j)|$. By the Reciprocity Theorem (Theorem 3.3), R is a symmetric matrix. Let $x_1 \in \mathbb{R}^m$ be a unit ℓ_1 normed vector that achieves the maximum dilation in the ℓ_1 norm when multiplied by R, i. e. $\|x_1\|_1 = 1$ and $\|Rx_1\|_1 = \|R\|_{1\to 1}$. We may assume without loss of generality that all coordinates of x_1 are non-negative, because if we replace x_1 by the vector x_1' obtained by taking absolute values of the coordinates of x_1, $\|Rx_1'\| \geq \|Rx_1\|$. Let $u_1, \dots, u_m$ be the standard basis. We note that

$$\begin{aligned} \|R\|_{1\to 1} &= \sum_{i=1}^{m} \|h_{e_i}\|_1 x_{1i} \\ &\leq \max_i \|h_{e_i}\|_1 \\ &= \max_i \|Ru_i\|_1 \\ &\leq \|R\|_{1\to 1}. \end{aligned}$$

Therefore $\|R\|_{1\to 1} = \max_i \|h_{e_i}\|_1$. Since R is a symmetric matrix whose entries are non-negative,

$$\sup_{\|x\|_\infty = 1} \|Rx\|_\infty = \|R(u_1 + \cdots + u_m)\|_\infty.$$

Therefore,

$$\begin{aligned} \|R\|_{\infty \to \infty} &= \|R(u_1 + \cdots + u_m)\|_\infty \\ &= \max_i \|h_{e_i}\|_1 \\ &= \|R\|_{1\to 1}. \end{aligned}$$

By Lemma 4.2 and the Riesz-Thorin theorem (Theorem 3.1), we conclude that

$$\begin{aligned} \sup_{D_p \in \bar{\mathcal{D}}_p} \frac{\|t_{D_p \cdot h}\|_p}{\|D_p\|_p} &\leq \|R\|_{p\to p} \\ &\leq \max(\|R\|_{1\to 1}, \|R\|_{\infty\to\infty}) \\ &= \|R\|_{1\to 1} (\text{By Lemma 4.2}) \\ &= \sup_{D_1 \in \bar{\mathcal{D}}_1} \frac{\|t_{D_1 \cdot h}\|_1}{\|D_1\|_1} \end{aligned}$$

which establishes Inequality 4.2 and thereby completes the proof.

Remark 1. *One can repeat the above argument using Mityagin's theorem (Theorem 3.2) instead of the Riesz-Thorin theorem, and prove the stronger statement that for any norm that transforms $\mathbb{R}^m$ into a symmetric unconditional Banach space X with respect to the standard basis,*

$$\kappa_X(\text{Harmonic}) \leq \kappa_1(\text{Harmonic}),$$

where $\kappa_X(\text{Harmonic})$ is the competitive ratio of Harmonic *with respect to the norm of X.*

4.1 Bounding π(Harmonic) by hitting and mixing times

Proposition 4.2. *For any vertices $i, j \in V$,*

$$\|h_{ij}\|_1 \leq 8T_{mix}(\frac{1}{4}).$$

Proof. [Proof of Proposition 4.2] It is always possible to find a potential ϕ_{ij} such that $\Delta\phi_{ij} = \delta_j - \delta_i$, and

$$\begin{aligned} \tilde{\pi}(\{u|\phi_{ij}(u) \leq 0\}) &\geq \frac{1}{2} \\ \tilde{\pi}(\{v|\phi_{ij}(v) \geq 0\}) &\geq \frac{1}{2}, \end{aligned}$$

because adding an arbitrary constant does not change the gradient of a potential. Let ϕ_{ij} satisfy the above conditions.

Recall that H_S be the (hitting) time taken for a random walk starting at v to hit the set S (Definition 4). Let $\mathbb{E}^v[\cdot]$ denote expectations assuming $W_0 = v$.

Lemma 4.3. *Suppose $\Delta\phi_{ij} = \delta_j - \delta_i$. Then,*

$$\|\nabla\phi_{ij}\|_1 \leq \sum_v \deg(v)|\phi_{ij}(v)|$$

Proof. [Proof of Lemma 4.3]

$$\begin{aligned} \|\nabla\phi_{ij}\|_1 &= \sum_{(u,v)\in E} |\phi_{ij}(u) - \phi_{ij}(v)| \\ &\leq \sum_{(u,v)\in E} (|\phi_{ij}(u)| + |\phi_{ij}(v)|) \\ &= \sum_v \deg(v)|\phi_{ij}(v)| \end{aligned}$$

Lemma 4.4. *Let $\Delta\phi_{ij} = \delta_j - \delta_i$, $S_\leq := \{v|\phi_{ij}(v) \leq 0\}$ and $S_\geq := \{v|\phi_{ij}(v) \geq 0\}$. Then,*

$$\sum_{v\in V} \deg(v)|\phi_{ij}(v)| \leq \mathbb{E}H^i_{S_\leq} + \mathbb{E}H^j_{S_\geq}.$$

Proof. [Proof of Lemma 4.4] Let $W_0, W_1, \dots$ be a random walk on G starting at i and ending the first time it hits $S_\leq$. Given $v \in V$ and a subset S of V, let $N^i_S(v)$ be the number of times the walk *exits* v until hitting $S_\leq$, and $\psi(v) := \frac{\mathbb{E}N^i_{S_\leq}(v)}{\deg(v)}$. Note that $\psi(u) = 0$ for all $u \in S_\leq$. We make the following claim.

CLAIM 1.

$$\begin{aligned}\Delta\psi(i) &= -1\\ \Delta\psi(u) &= 0 \quad \textit{if} \;\; u \in V \setminus \{S_{\leq} \cup \{i\}\}.\end{aligned}$$

Proof. [Proof of Claim 1] To see why this is true, let $E(t, v)$ be the event that $W_t = v$. For any vertex u, let $\star(u)$ denote the set of vertices adjacent to u. We see that for $1 \leq t \leq H^t_{S_{\leq}}$ and $u \in V \setminus \{S_{\leq} \cup \{i\}\}$,

$$\mathbb{P}[E(t, v)] = \sum_{u \in \star(v)} \frac{\mathbb{P}[E(t-1, u)]}{\deg(u)}.$$

Summing up over time, this implies $\mathbb{E}[N^i_{S_{\leq}}(v)] = \sum_{u\in\star(v)} \frac{\mathbb{E}[N^i_{S_{\leq}}(u)]}{\deg(u)}$. This translates to $\psi(v) = \sum_{u\in\star(v)} \frac{\psi(u)}{\deg(v)}$. When $v = i$, a similar computation yields

$$\psi(i) = 1 + \sum_{u\in\star(i)} \frac{\psi(u)}{\deg(i)},$$

proving the claim.

It follows that $\Delta(\psi - \phi)$ is 0 on *all* of $V \setminus S_{\leq}$. This implies that the maximum of $\phi - \psi$ cannot be achieved on $V \setminus S_{\leq}$ (Maximum principle for Harmonic functions). $\phi - \psi$ is ≤ 0 on $S_{\leq}$, therefore $\psi - \phi$ is a non-negative function. It follows that

$$\sum_{v\in S_{\geq}} \deg(v)\phi_{ij}(v) \leq \mathbb{E}^i\left[H_{S_{\leq}}\right].$$

An identical argument applied to $-\phi_{ij}$ instead of ϕ_{ij} gives us the following.

$$\sum_{v\in S_{\leq}} \deg(v)(-\phi_{ij}(v)) \leq \mathbb{E}^j\left[H_{S_{\geq}}\right].$$

Together, the last two inequalities complete the proof.

LEMMA 4.5. *Let S be a subset of V whose stationary measure is greater or equal to $\frac{1}{2}$. Let $v \in V \setminus S$. Then*

$$\mathbb{E}^v[H_S] \leq 4T_{mix}(1/4).$$

Proof. [Proof of Lemma 4.5] Let $W_0, W_1, \ldots$ be a random walk on G starting at v. Recall that from Definition 5,

$$T_{mix}(\epsilon) = \sup_{v\in V} \inf\left\{T \middle| \forall_{t\geq T} \|\tilde{\pi} - \rho_v^{(t)}\|_1 < \epsilon\right\}.$$

Let $\tau := T_{mix}(1/4)$ and $\mathbb{P} = \mathbb{P}^v$.

$$\begin{aligned}\mathbb{P}[H_S \leq \tau] &\geq \mathbb{P}[W_\tau \in S]\\ &= \tilde{\pi}(S) - (\tilde{\pi}(S) - \sum_{u\in S} \mathbb{P}[W_\tau = u])\\ &\geq \frac{1}{2} - |\tilde{\pi}(S) - \sum_{u\in S}\mathbb{P}[W_\tau = u]|\\ &\geq \frac{1}{2} - |\sum_{u\in S}(\tilde{\pi}(u) - \mathbb{P}[W_\tau = u])| \geq \frac{1}{4}.\end{aligned}$$

In order to get a bound on the expected hitting time from this bound on the hitting probability, we observe that the distribution of hitting times has an exponential tail. More precisely, using the Markovian property of the random walk,

$\mathbb{P}[H_S > k\tau]$ is less or equal to

$$\mathbb{P}\left[X_\tau \notin S\right] \prod_{i=1}^{k-1} \sup_{u\notin S} \mathbb{P}\left[X_{(i+1)\tau} \notin S \middle| X_{i\tau} = u\right] \leq \frac{3^k}{4^k}.$$

Finally,

$$\mathbb{E}^v[H_S] \leq \tau\left(\sum_{i=0}^{\infty} \mathbb{P}\left[H_S^v > i\tau\right]\right) \leq 4\tau.$$

PROPOSITION 4.3. *For any edge e, $\|h_e\|_1 \leq \sqrt{m}$.*

Proof. [Proof of Proposition 4.3] If $e = (u, v)$, h_e is by Thompson's principle ([4]) the minimizer of $\|f\|_2$ among all flows f for which $\operatorname{div} f = \delta_u - \delta_v$. Since u and v are adjacent, $\|f\|_2 = 1$ if f is the "shortest-path" flow defined by

$$f(e') = \begin{cases} 1 & \text{if } e' = e\\ 0 & \text{otherwise.}\end{cases}$$

Therefore $\|h_e\|_2 \leq 1$. This implies that $\|h_e\|_1 \leq \sqrt{m}$, since for any vector $x \in \mathbb{R}^m$, $\|x\|_1 \leq \sqrt{m}\|x\|_2$.

Proof. [Proof of Theorem 2.3]

Let G be a graph with $n = m - \lfloor\sqrt{m-1}\rfloor + 1$ vertices and m edges constructed as follows (see Figure 1). Let vertices labeled 1 and 2 be joined by an edge e_1, and also be connected by $r := \lfloor\sqrt{m-1}\rfloor$ vertex disjoint paths of length h. We make the remaining vertices and edges belong to a path from vertex 3 to 2, such that there is no path from any of these vertices to 1 which does not contain 2. We will fix a specific set of demands $D = \{d_{ij}\}$; namely a unit demand from 1 to 2, and 0 otherwise. Suppose that an algorithm $\mathcal{A}$ uses a flow denoted a_{12} to achieve this, where $a_{12}(e_1) = \alpha \in [0, 1]$. Then,

$$\|a_{ij}\|_1 \geq \alpha + (1 - \alpha)r,$$

and

$$\|a_{ij}\|_\infty \geq \alpha.$$

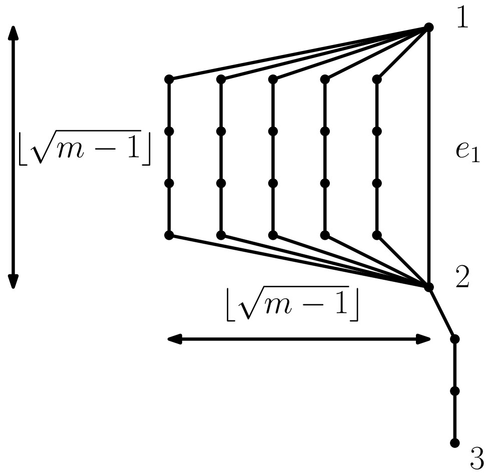

Figure 1: A graph on which the performance index $\pi(\mathcal{A})$ of any oblivious algorithm $\mathcal{A}$ is $\geq \frac{\lfloor\sqrt{m-1}\rfloor}{2}$

On the other hand a flow $\textsc{Opt}_1$ that uses only e_1 incurs an ℓ_1 norm of 1. A flow $\textsc{Opt}_\infty$ that uses all the $r+1$ edge disjoint paths from 1 to 2 equally incurs an ℓ_∞ cost of $\frac{1}{r+1}$.

$$\frac{\|a_{12}\|_1}{\|\textsc{Opt}_1\|_1} \geq (1-\alpha)r + \alpha,$$

and

$$\frac{\|a_{12}\|_\infty}{\|\textsc{Opt}_\infty\|_\infty} \geq \alpha(r+1).$$

Therefore,

$$\max(\kappa_1(\mathcal{A}), \kappa_\infty(\mathcal{A})) \geq \frac{\kappa_1(\mathcal{A}) + \kappa_\infty(\mathcal{A})}{2} \geq \frac{\lfloor\sqrt{m-1}\rfloor}{2}.$$

5 Random walks and π(Harmonic)

5.1 Random walk on the torus

Proof. [Proof of Theorem 2.2] We consider the example of simple random walk $W_0, W_1, \ldots$ on the 2-dimensional torus

$$\mathbb{T}_N = \{(x_1, x_2) \in \mathbb{Z}^2 : x_j \in \{0, 1, \ldots, N-1\},$$

with the usual periodic boundary condition. Note that $\mathbb{T}_N$ contains N^2 vertices and $2N^2$ edges. The local central limit theorem implies that T_{mix} is $\Theta(N^2)$.

Let u denote a nearest neighbor of the origin, and let $\phi = \phi_{0u}$ be the corresponding potential with the additive constant chosen so that $\phi(0) = -\phi(u)$. In other words, for every vertex x,

$$\phi(x) = \frac{b}{2}\,\mathbb{P}^x\{W(\overline{H}_{0u}) = 0\} - \frac{b}{2}\,\mathbb{P}^x\{W(\overline{H}_{0u}) = u\}$$

$$= b\,\mathbb{P}^x\{W(\overline{H}_{0u}) = 0\} - \frac{b}{2},$$

where

$$\frac{1}{b} = \mathbb{P}^u\{W(H_{0u}) = 0\} = \mathbb{P}^0\{W(H_{0u}) = u\} \geq \frac{1}{4}.$$

Proposition 5.1. *As $N \to \infty$,*

$$\|\nabla\phi\|_1 = O(\log N).$$

Proof. [Proof of Proposition 5.1] We will prove the stronger statement, that if $e = (x, y)$ is an edge,

$$|\nabla\phi(e)| \leq c\,|x|^{-2}, \tag{5.5}$$

where the distance $|\cdot|$ is taken on the torus $\mathbb{T}_N$. In order to prove this, it suffices to prove that

$$|\phi(x)| \leq c\,|x|^{-1}. \tag{5.6}$$

Indeed, if this holds then $|\phi(x)| \leq 2c\,|x|^{-1}$ in the disk of radius $|x|/2$ about x and then standard difference estimates for discrete harmonic functions (see, e.g., [8, Theorem 1.7.1]) imply (5.5). Since $b \leq 2d$, (5.6) follows from

$$\mathbb{P}^x\{W(H_{0,u}) = 0\} = \mathbb{P}^x\{W(H_{0,u}) = u\} + O(|x|^{-1}),$$

which we now prove. Let $C = C_x = \{v : |v| < |x|/2\}$, $\partial C = \{y \notin C : |v - y| = 1 \text{ for some } v \in C\}$, $S = \partial C \cup \{0, u\}$. By focusing at the last visit to ∂C before reaching $\{0, u\}$, we can see that

$$\mathbb{P}^x\{W(H_{0,u}) = 0\} = \sum_{y \in \partial C} G(x, y)\,\mathbb{P}^y\{W(H_S) = 0\}$$

$$= \sum_{y \in \partial C} G(x, y)\,\mathbb{P}^0\{W(H_S) = y\},$$

where $G = G_{0,u}$ is the Green's function for the simple random walk killed when it reaches $\{0, u\}$. The second inequality uses a symmetry argument. A similar expression holds for $\mathbb{P}^x\{W(H_{0,u}) = u\}$, so it suffices to show for $y \in \partial C$,

$$\mathbb{P}^0\{W(H_S) = y\} = \mathbb{P}^u\{W(H_S) = y\}\,\left[1 + O(|x|^{-1})\right].$$

This estimate follows from the following estimates:

$$\mathbb{P}^0\{W(H_{\partial C}) = y\} \leq \frac{c}{|x|}\mathbb{P}^v\{W(H_S) \in \partial C\} \leq \frac{c}{\log|x|}, \quad v \in \{0,u\},$$

$$\mathbb{P}^v\{W(H_S) = y \mid W(H_S) \in \partial C\} = \mathbb{P}^0\{W(H_{\partial C}) = y\}\left[1 + O\left(\frac{\log|x|}{|x|}\right)\right],$$

for $v \in \{0,u\}, y \in \partial C$. For proofs of these, see Theorem 1.6.5, Lemma 1.7.4, and Theorem 2.1.3, respectively, of [8].

The last proposition can be extended to the skewed torus

$$\mathbb{T}(N,M) = \{(x_1,x_2) : x_1 = 0,\ldots,N-1; x_2 = 0,\ldots,M-1\},$$

with $N \leq M$ to show that in this case

$$\|\nabla\phi\| \leq c \log N,$$

where c is independent of M. The proof is similar, but one gives a coupling argument to show that for $|x| \geq N$,

$$|\phi(x)| \leq c\,N^{-1}\,e^{-\beta|x|/N},$$

for some constants c,β. The difference estimate on the disk of radius $N/2$ about x, then yields

$$|\nabla\phi(e)| \leq c\,N^{-2}\,e^{-\beta|x|/N},$$

for for any edge e that has x as an endpoint.

5.2 A Bound on the Spectral Gap using flows of Harmonic

A well known method method due to Diaconis-Stroock [3] and Sinclair [15], of bounding the spectral gap of a reversible Markov chain involves the construction of canonical paths from each node to every other. One may use flows instead of paths, and derive better mixing bounds, as was the case in the work of Morris-Sinclair [10] on sampling knapsack solutions. Sinclair suggested a natural way of constructing canonical flows using random walks in [15]. This scheme gives a bound of $O(\tau^2)$, if τ is the true mixing time. In this section, we observe that for a random walk on a graph, the flows of Harmonic provide a certificate for rapid mixing as well, and give the same $O(\tau^2)$ bound on the mixing time. Let the stationary distribution be denoted $\tilde{\pi}$. $\tilde{\pi(i)} = \frac{\deg(i)}{2m}$. Let the *capacity* of an edge $e = (u,v)$, denoted $Q(e)$ be defined to be $\tilde{\pi}(u)p_{uv}$, where $p_{uv} = \frac{1}{\deg(u)}$ is the transition probability from u to v. Let the transition matrix be denoted P. In our setting of (unweighted) random walks, $\tilde{\pi(i)} = \frac{\deg(i)}{2m}$ and for each edge e, $Q(e)$ is $\frac{1}{2m}$.

Definition 7. *Let D be the demand $\langle d_{ij} : i,j \in V\rangle$, where $d_{ij} = \tilde{\pi}(i)\tilde{\pi}(j)$. Given a multi-flow f, where $f \searrow D$, let $\rho(f)$ denote the maximum load on any edge divided by its capacity i. e. $\rho(f) = \frac{\|t_f\|_\infty}{Q(e)}$.*

Theorem 5.1. (Sinclair) *Let $f \searrow D$ as described above. Then, the second eigenvalue λ_1 of the transition matrix P satisfies*

$$\lambda_1 \leq 1 - \frac{1}{8\rho^2}.$$

Let us denote $T_{mix}(1/4)$ by τ. Let $h \cdot D$ be the multi-flow obtained by re-scaling h so that it meets D. Then, we have the following proposition.

Proposition 5.2.

$$\frac{1}{2\sqrt{2(1-\lambda_1)}} \leq \rho(h \cdot D) \leq 16\tau.$$

Proof. [Proof of Proposition 5.2] The lower bound on $\rho(h \cdot D)$ follows from Theorem 5.1. We proceed to show the upper bound. In the case of a random walk on a graph with m edges, for every edge e, $Q(e) = \frac{1}{2m}$. Therefore

$$\rho(D \cdot h) = \frac{\|t_{D\cdot h}\|_\infty}{Q(e)} = 2m\|t_{D\cdot h}\|_\infty.$$

Given an edge $e = (u,w)$, let $\phi_e = \phi_{uw}$ be a potential such that $h_{uw} = \nabla\phi_{uw}$. For convenience, for any i, let h_{ii} be defined to be the flow that is zero on all edges. Let $e = (u,w)$ be the edge that carries the maximum load. Then,

$$\begin{aligned}
\rho(D \cdot h) &= \sum_{i,j} \frac{|\deg(i)\deg(j)h_{ij}(e)|}{2m} \\
&= \sum_{i,j} \frac{\deg(i)\deg(j)}{2m}|\phi_{ij}(u) - \phi_{ij}(w)| \\
&= \sum_{i,j} \frac{\deg(i)\deg(j)}{2m}|\phi_e(i) - \phi_e(j)| \quad \text{(Thm 3.3)} \\
&\leq \sum_{i,j} \frac{\deg(i)\deg(j)}{2m}(|\phi_e(i)| + |\phi_e(j)|) \\
&= 2\sum_i \deg(i)|\phi_e(i)|.
\end{aligned}$$

Adding an appropriate constant to ϕ_e is necessary, we may assume that

$$\tilde{\pi}(\{u|\phi_{ij}(u) \leq 0\}) \geq \frac{1}{2}$$

$$\tilde{\pi}(\{v|\phi_{ij}(v) \geq 0\}) \geq \frac{1}{2}.$$

Then, Lemma 4.4 and Lemma 4.5 together imply that

$$2\sum_{i}\deg(i)|\phi_e(i)| \leq 16\tau,$$

completing the proof.

References

[1] B. Awerbuch, Y. Azar, E. F. Grove, M.-Y. Kao, P. Krishnan, and J. S. Vitter, *Load balancing in the L_p norm.*, in Proceedings of the 36th IEEE Symposium on Foundations of Computer Science (FOCS), 1995, pp. 383–391.

[2] Y. Bartal and S. Leonardi, *On-line routing in all optimal networks,* Theoretical Computer Science, 1997.

[3] P. Diaconis and D. Stroock, *Geometric bounds for eigenvalues of Markov Chains*, Annals of Applied Probability, 1, (1991).

[4] P. G. Doyle and J. L. Snell, *Random Walks and Electric Networks*, Mathematical Association of America, 1984.

[5] C. Harrelson, K. Hildrum, and S. B. Rao, *A polynomial-time tree decomposition to minimize congestion*, in Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2003, pp. 34–43.

[6] M. Hajiaghayi, R. Kleinberg, T. Leighton, and H. Räcke *New lower bounds for oblivious routing in undirected graphs* In Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA 2006), pp. 918-927.

[7] P. Harsha, T. Hayes, H. Narayanan, H. Räcke, and J. Radhakrishnan, *Minimizing average latency in Oblivious routing*, (SODA), 2008.

[8] G. Lawler *Intersections of Random Walks* Birkhäuser, 1991.

[9] B.S. Mityagin *An interpolation theorem for modular spaces* , Mat. Sb. (N.S.) 66 (108) 1965 473482

[10] B. Morris and A. Sinclair, *Random Walks on Truncated Cubes and Sampling 0-1 knapsack solutions.* SIAM journal on computing 34 (2004), pp. 195-226

[11] H. Räcke, *Minimizing congestion in general networks*, in Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS), 2002, pp. 43–52.

[12] H. Räcke, *Optimal Hierarchical Decompositions for Congestion Minimization in Networks.* In Proc. of the 40th STOC, 2008.

[13] M. Riesz, *Sur les maxima des formes bilinéaires et sur les fonctionelles linéaires.* Acta Math. 49, 465 - 497 (1926).

[14] T. Roughgarden and É. Tardos, *How bad is selfish routing?*, Journal of the ACM, 49 (2002), pp. 236–259.

[15] A. Sinclair *Improved Bounds for Mixing Rates of Markov Chains and Multicommodity Flow.* Combinatorics, Probability and Computing 1 (1992), pp. 351-370

[16] G. O. Thorin, *Convexity theorems generalizing those of M. Riesz and Hadamard with some applications.* Comm. Sem. Math. Univ. Lund = Medd. Lunds Univ. Sem. 9, 1 - 58 (1948).

[17] P. Tetali, *Random walks and the effective resistance of networks* Journal of Theoretical Probability. Volume 4, Number 1, January, 1991

[18] L. G. Valiant and G. J. Brebner, *Universal schemes for parallel communication*, in Proceedings of the 13th ACM Symposium on Theory of Computing (STOC), 1981, pp. 263–277.