

Duality and evolving set bounds on mixing times

Ravi Montenegro ¹

November 20, 2007

¹Department of Mathematical Sciences, University of Massachusetts Lowell, Lowell, MA 01854, ravi.montenegro@uml.edu; partially supported by a VIGRE grant at the Georgia Institute of Technology.

Abstract

We sharpen the Evolving set methodology of Morris and Peres and extend it to study convergence in total variation, relative entropy, L^2 and other distances. Bounds in terms of a modified form of conductance are given which apply even for walks with no holding probability. These bounds are found to be strictly better than earlier Evolving set bounds, may be substantially better than conductance profile results derived via Spectral profile, drastically sharpen Blocking Conductance bounds if there are no bottlenecks at small sets, and give intuition into the workings of Canonical Path methods.

This paper is intended solely to develop theoretical underpinnings, and as such we focus on two points : proving the sharpest most general results we can, and showing the Evolving Set methods to be better than previous isoperimetric methods. In order to learn about Evolving Sets we recommend the relevant chapter in our book with Tetali [23], and of course the original paper of Morris and Peres [26]. To learn about some applications please see our paper on Cheeger Inequalities [21], that on Canonical Path bounds for non-lazy walks [18], our alternate interpretation of Morris' study of the Thorp shuffle [23, 25], Morris' paper on the Exclusion process [24], and the paper of Diaconis and Fill on the duality method [2].

Keywords : Mixing time, evolving sets, blocking conductance, spectral profile, conductance.

Contents

- 1 Introduction** **2**

- 2 Set bounds on distance and Mixing Times** **4**
 - 2.1 Duality and Evolving sets 4
 - 2.2 Evolving set bounds on distances 6
 - 2.3 Mixing times 8
 - 2.4 Continuous Time 11

- 3 Conductance and Modified Conductance** **12**
 - 3.1 Conductance 12
 - 3.2 Modified conductance 13
 - 3.3 An Inequality Prover 14
 - 3.4 Bounds on f -congestion $\mathcal{C}_f(A)$ 16
 - 3.5 Flow distribution 20
 - 3.6 The effect of re-scaling on mixing time 22
 - 3.7 Examples 26

- 4 A comparison to previous isoperimetric bounds** **31**
 - 4.1 Evolving Sets 31
 - 4.2 Spectral Profile 31
 - 4.3 Blocking Conductance 34
 - 4.4 Comparison and Canonical paths for Evolving Sets 39
 - 4.5 An aside into Blocking Conductance and comparison 46
 - 4.6 Spread and the discrete gradients 48

- Bibliography** **53**

Chapter 1

Introduction

An isoperimetric bound on mixing time uses a geometric quantity, such as conductance, to bound the rate of convergence of a Markov chain. Such bounds have played a key role in proving mixing time results, beginning with Jerrum and Sinclair's [11] proof that a random walk for approximating the permanent of a dense matrix converges in polynomial time. Their idea has been extended to apply to non-reversible non-lazy walks [16, 6], to continuous state spaces [15], to walks with low conductance on small sets [15], and to walks with high conductance on small sets [14].

Three recent papers have built on the Average Conductance idea of Lovász and Kannan [14]. Morris and Peres [26] develop the Evolving Set methodology to show very strong results in terms of L^2 distance. Kannan, Lovász and Montenegro [13] show similar results for total variation distance of a reversible, lazy walk through the method of Blocking Conductance. Finally, Goel, Montenegro, and Tetali [9] use the notion of Spectral Profile to extend an approach of Fill [6] and bound L^2 mixing of finite Markov chains. Each of these were shown by very different methods: by using a duality based approach, by considering the n -step average distribution, and by direct examination of the drop in variance, respectively.

The goal of this paper is to develop a general framework under which these isoperimetric results are unified as much as possible. This will be done by strengthening the Evolving Set methodology. Our improved argument leads to bounds on any convex notion of distance: including total variation, relative entropy, L^2 , Hellinger, and Wasserstein distances. These are the first isoperimetric bounds on most of these distances, and even when past bounds are known these are the first which are sharp. For each of these distances we can also derive bounds in terms of an extension of the conductance method, known as modified conductance, which is consistent with past bounds when applied to lazy walks but which also applies in the setting of walks with no holding probability.

How do our new Evolving Set results compare to previous isoperimetric bounds? We find that our new L^2 mixing bound is slightly better than earlier Evolving Set results, our conductance bounds on L^2 mixing may be substantially better than those derived from Spectral Profile bounds, and our mixing bounds are significantly sharper than those of Blocking Conductance except when the worst bottleneck is at a small set. Moreover, our results explain the curious existence of three total variation mixing bounds in the Blocking Conductance paper [13]. We find these are in fact total variation extensions of a bound on L^2 mixing, a bound on relative entropy mixing, and a direct bound on total variation mixing. An Evolving set approach to canonical paths also suggests that previous forms in terms of edge-expansion, vertex-expansion, or path lengths can all be considered to be bounds of the form (edge-expansion)*(vertex-expansion).

This paper is focused on developing a rich theoretical framework, and comparing it to past methods. As such it is not so much as a text on applying Evolving Sets, as a text developing theory. Some examples are, however, included at the end in the Examples section. The interested reader can find additional theoretical developments in [21], where we show a version of Cheeger's inequality which bounds (complex-valued) eigenvalues of non-reversible chains, a version to bound the smallest eigenvalue of a reversible chain, and we also sharpen Cheeger inequalities of Jerrum and Sinclair, Alon, and Stoyanov for bounding the spectral gap in terms of isoperimetric measures of edge and/or vertex expansion of a non-reversible walk. In [18] we develop a canonical path bound for non-reversible non-lazy walks, and use this to extend past results on mixing times of reversible walks on Cayley graphs into the general setting. Finally, together with Tetali [23] we substantially improve on mixing time bounds of Morris for the Thorp shuffle [25], by use of a conductance-profile bound based on ideas developed in this paper for walks with no holding probability.

The paper proceeds as follows. In Section 2 we introduce the notion of Evolving sets, and use this to show isoperimetric bounds on distances and mixing times. This is followed in Section 3 by conductance and modified conductance, an extension of conductance to non-lazy walks. These new results are compared to previous isoperimetric methods in Section 4.

Chapter 2

Set bounds on distance and Mixing Times

In this section the main development of this paper is given: isoperimetric methods for bounding several notions of distance and mixing time. The arguments are based on the evolving set process of Morris and Peres [26] which was also described in the context of duality by Diaconis and Fill [2].

A little notation is required. Let \mathbf{P} be a finite irreducible Markov kernel on state space V with stationary distribution π , that is, \mathbf{P} is a $|V| \times |V|$ matrix with entries in $[0, 1]$, row sums are one, V is connected under \mathbf{P} (i.e. $\forall x, y \in V \exists t : \mathbf{P}^t(x, y) > 0$), and π is a distribution on V with $\pi\mathbf{P} = \pi$. The time-reversal \mathbf{P}^* is given by $\mathbf{P}^*(x, y) = \frac{\pi(y)\mathbf{P}(y, x)}{\pi(x)}$ and is a Markov chain with stationary distribution π as well. If $A, B \subset V$ the ergodic flow from A to B is given by $\mathbf{Q}(A, B) = \sum_{x \in A, y \in B} \pi(x)\mathbf{P}(x, y)$. Given initial distribution σ , the n -step discrete time distribution is given by $\sigma\mathbf{P}^n$, and if the walk is aperiodic then $\sigma\mathbf{P}^n \xrightarrow{n \rightarrow \infty} \pi$.

2.1 Duality and Evolving sets

In order to relate a property of sets (conductance) to a property of the original walk (mixing time) we construct a walk on sets that is a dual to the original Markov chain. Given a Markov chain on V with transition matrix \mathbf{P} , a *dual process* consists of a walk $\hat{\mathbf{P}}$ on some state space V and a *link*, or transition matrix, Λ from V to V such that

$$\mathbf{P}\Lambda = \Lambda\hat{\mathbf{P}}.$$

In particular, $\mathbf{P}^n\Lambda = \Lambda\hat{\mathbf{P}}^n$ and so the evolution of \mathbf{P}^n and $\hat{\mathbf{P}}^n$ will be closely related. This relation is given visually by Figure 2.1.

The projection $\Lambda(S, y) = \frac{\pi(y)}{\pi(S)} 1_S(y)$ is a natural candidate to link a walk on sets to a walk on states. Diaconis and Fill [2] have shown that for certain classes of Markov chains that the walk $\hat{\mathbf{K}}$ below is the unique dual process with link Λ , so this is the walk on sets that should be considered. We use notation of Morris and Peres [26].

Definition 2.1.1. Given set $A \subset V$ a step of the *evolving set process* is given by choosing $u \in [0, 1]$ uniformly at random, and transitioning to the set $A_u = \{y \in V : \mathbf{Q}(A, y) \geq u\pi(y)\}$. The walk is denoted by $S_0, S_1, S_2, \dots, S_n$, with transition kernel $\mathbf{K}^n(A, S) = \text{Prob}(S_n = S | S_0 = A)$.




Figure 2.1: The dual walk $P_D = \hat{P}$ projects onto the original chain P .

Definition 2.1.2. The *Doob transform* of the Evolving set process is the Markov chain \hat{K} on sets with transition probabilities

$$\hat{K}(S, S') = \frac{\pi(S')}{\pi(S)} \mathbf{K}(S, S').$$

The n -step transition probabilities are $\hat{K}^n(S, S') = \frac{\pi(S')}{\pi(S)} \mathbf{K}^n(S, S')$.

The Doob transform produces a Markov chain because of a Martingale property.

Lemma 2.1.3. *If $A \subset V$ then*

$$\int_0^1 \pi(A_u) du = \pi(A).$$

Proof.

$$\int_0^1 \pi(A_u) du = \sum_{y \in V} \pi(y) \text{Prob}(y \in A_u) = \sum_{y \in V} \pi(y) \frac{\mathbf{Q}(A, y)}{\pi(y)} = \pi(A).$$

□

The walk \hat{K} is a dual process of P .

Lemma 2.1.4. *If $S \subset V$, $y \in V$ and $\Lambda(S, y) = \frac{\pi(y)}{\pi(S)} \mathbf{1}_S(y)$ is the projection linkage, then*

$$P\Lambda(S, y) = \Lambda\hat{K}(S, y).$$

Proof.

$$\begin{aligned} P\Lambda(S, y) &= \sum_{z \in S} \frac{\pi(z)}{\pi(S)} P(z, y) = \frac{\mathbf{Q}(S, y)}{\pi(S)} \\ \Lambda\hat{K}(S, y) &= \sum_{S' \ni y} \hat{K}(S, S') \frac{\pi(y)}{\pi(S')} = \frac{\pi(y)}{\pi(S)} \sum_{S' \ni y} \mathbf{K}(S, S') = \frac{\mathbf{Q}(S, y)}{\pi(S)} \end{aligned}$$

The final equality is because $\sum_{S' \ni y} \mathbf{K}(S, S') = \text{Prob}(y \in S') = \mathbf{Q}(S, y)/\pi(y)$.

□

With duality it becomes easy to write the n step density in terms of the walk \hat{K} .

Lemma 2.1.5. *Let $\hat{\mathbb{E}}_n$ denote expectation under \hat{K}^n . If $x \in V$ and $S_0 = \{x\}$ then*

$$P^n(x, y) = \hat{\mathbb{E}}_n \pi_{S_n}(y),$$

where $\pi_S(y) = \frac{\mathbf{1}_S(y)\pi(y)}{\pi(S)}$ denotes the probability distribution induced on set S by π .

Proof.

$$\mathbf{P}^n(x, y) = \mathbf{P}^n \Lambda(\{x\}, y) = \Lambda \hat{\mathbf{K}}^n(\{x\}, y) = \hat{\mathbb{E}}_n \pi_{S_n}(y)$$

The final equality is because $\Lambda(S, y) = \pi_S(y)$. □

2.2 Evolving set bounds on distances

It is now a short hop from Lemma 2.1.5 to a bound on mixing times. First, however, note that if a distance $dist(\mu, \pi)$ is convex in μ (i.e. $dist(\alpha\mu_1 + (1-\alpha)\mu_2, \pi) \leq \alpha dist(\mu_1, \pi) + (1-\alpha)dist(\mu_2, \pi)$), then for any distribution σ

$$dist(\sigma \mathbf{P}^n, \pi) = dist\left(\sum_{x \in V} \sigma(x) \mathbf{P}^n(x, \cdot), \pi\right) \leq \sum_{x \in V} \sigma(x) dist(\mathbf{P}^n(x, \cdot), \pi) \leq \max_{x \in V} dist(\mathbf{P}^n(x, \cdot), \pi).$$

In this case distance is maximized when the initial distribution is a point mass, i.e. $\sigma(y) = \delta_{y=x}$ for some $x \in V$. Given the preceding lemmas it is easy to show an evolving set bound for all convex distances.

Lemma 2.2.1. *Consider a finite Markov chain with stationary distribution π . Any distance $dist(\mu, \pi)$ which is convex in μ satisfies*

$$dist(\mathbf{P}^n(x, \cdot), \pi) \leq \hat{\mathbb{E}}_n dist(\pi_{S_n}, \pi)$$

whenever $x \in V$ and $S_0 = \{x\}$.

Proof. By Lemma 2.1.5 and convexity,

$$dist(\mathbf{P}^n(x, \cdot), \pi) = dist(\hat{\mathbb{E}}_n \pi_{S_n}, \pi) \leq \hat{\mathbb{E}}_n dist(\pi_{S_n}, \pi).$$

□

Many distances are used in studying mixing times. These include:

- *Separation distance:* $s(\mu, \pi) = \max_{x \in V} 1 - \frac{\mu(x)}{\pi(x)}$
- *Total variation distance:* $\|\mu - \pi\|_{TV} = \frac{1}{2} \sum_{x \in V} |\mu(x) - \pi(x)|$
- *Relative Entropy:* $D(\mu \|\pi) = \sum_{x \in V} \mu(x) \log \frac{\mu(x)}{\pi(x)}$
- *L^2 distance:* $\left\| \frac{\mu}{\pi} - 1 \right\|_{2, \pi} = \sqrt{\sum_{x \in V} \pi(x) \left| \frac{\mu(x)}{\pi(x)} - 1 \right|^2}$
- *Relative Pointwise distance (L^∞):* $\left\| \frac{\mu}{\pi} - 1 \right\|_\infty = \max_{x \in V} \left| \frac{\mu(x)}{\pi(x)} - 1 \right|$
- *Hellinger distance:* $H(\mu, \pi) = \sum_{y \in V} \left(\sqrt{\frac{\mu(y)}{\pi(y)}} - 1 \right)^2 \pi(y)$

- *Wasserstein distance* $W_p(\mu, \pi)$: Given metric $d : V \times V \rightarrow \mathbb{R}^+$, let

$$W_p^p(\mu, \pi) = \sup_{\substack{f, g: V \rightarrow \mathbb{R}, \\ \forall y, z \in V: f(y) + g(z) \leq d(y, z)^p}} \mathbb{E}_\mu f + \mathbb{E}_\pi g$$

Each of these distances can be bounded easily with Lemma 2.2.1.

Theorem 2.2.2. *Given a finite, ergodic Markov chain, $x, y \in V$ and $S_0 = \{x\}$, then*

$$\begin{aligned} s(\mathbf{P}^n(x, \cdot), \pi) &\leq \text{Prob}_{\hat{\kappa}^n}(S_n \neq V) \\ \|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} &\leq \hat{\mathbb{E}}_n(1 - \pi(S_n)) \\ D(\mathbf{P}^n(x, \cdot) \|\pi) &\leq \hat{\mathbb{E}}_n \log \frac{1}{\pi(S_n)} \\ \|\mathbf{P}^n(x, \cdot) - \pi\|_{2, \pi} &\leq \hat{\mathbb{E}}_n \sqrt{\frac{1 - \pi(S_n)}{\pi(S_n)}} \\ \left| \frac{\mathbf{P}^n(x, y)}{\pi(y)} - 1 \right| &\leq \max \left\{ \frac{1 - \pi(y)}{\pi(y)}, 1 \right\} \text{Prob}_{\hat{\kappa}^n}(S_n \neq V) \\ H(\mathbf{P}^n(x, \cdot), \pi) &\leq 2\hat{\mathbb{E}}_n(1 - \sqrt{\pi(S_n)}) \\ W_p(\mathbf{P}^n(x, \cdot), \pi) &\leq \sqrt[p]{\hat{\mathbb{E}}_n W_p^p(\pi(S_n), \pi)} \end{aligned}$$

Most of these are immediate from the lemma and computation of $\text{dist}(\pi_S, \pi)$. For instance, in the total variation case $\|\pi_S - \pi\|_{TV} = 1 - \pi(S)$.

A few cases are worth mentioning further. The relative pointwise bound is because $\text{dist}(\mu, \pi) = \left| \frac{\mu(y)}{\pi(y)} - 1 \right|$ is convex, with

$$\text{dist}(\pi_S, \pi) = \left| \frac{\pi_S(y)}{\pi(y)} - 1 \right| = \mathbf{1}_S(y) \frac{1 - \pi(S)}{\pi(S)} + \mathbf{1}_{S^c}(y) \leq \max \left\{ \frac{1 - \pi(y)}{\pi(y)}, 1 \right\} \delta_{S \neq V}.$$

The Hellinger distance is a special case of $\text{dist}(\mu, \pi) = \mathcal{L}_\pi \left(\frac{\mu}{\pi} \right)$ for a convex functional $\mathcal{L}_\pi : (\mathbb{R}^+)^V \rightarrow \mathbb{R}$. Wasserstein distance is a case of $\mathcal{L}_\pi(f) = \sup_{h \in H} \sum_{y \in V} h(y) f(y) \pi(y)$ for some class of functions H , by rewriting as

$$W_p^p(\mathbf{P}^n(x, \cdot), \pi) = \sup_{\substack{f, g: V \rightarrow \mathbb{R}, \\ \forall y, z \in V: f(y) + g(z) \leq d(y, z)^p}} \sum_{y \in V} (f(y) + \mathbb{E}_\pi g) \left(\frac{\mu(y)}{\pi(y)} \right) \pi(y).$$

The Wasserstein distance W_p^p is just the total variation distance when $d(y, z) = \delta_{y \neq z}$. It is easily checked that $W_p^p(\pi_S, \pi) = 1 - \pi(S)$ in this case, and so

$$\|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} = W_p^p(\mathbf{P}^n(x, \cdot), \pi) \leq \hat{\mathbb{E}}_n W_p^p(\pi_{S_n}, \pi) = \hat{\mathbb{E}}_n(1 - \pi(S_n)),$$

which shows the Wasserstein bound generalizes the total variation bound.

Remark 2.2.3. *When the initial distribution σ is not a point mass then S_0 should be chosen from a distribution. Set $\sigma_0 = \sigma$. Inductively define $A_i = \{x : \sigma_i(x) > 0\}$, let $\text{Prob}(S_0 = A_i) =$*

$\pi(A_i) \min_{x \in A_i} \frac{\sigma_i(x)}{\pi(x)}$, and $\sigma_{i+1}(x) = \sigma_i(x) - \mathbf{1}_{A_i}(x) \frac{\pi(x)}{\pi(A_i)} \text{Prob}(S_0 = A_i)$. Then Lemma 2.1.5 generalizes to

$$\frac{\sigma \mathbf{P}^n(y)}{\pi(y)} = \sum_{S \subset V} \text{Prob}(S_0 = S) \frac{\text{Prob}(y \in S_n | S_0 = S)}{\pi(S)} = \hat{\mathbb{E}}_n \frac{\mathbf{1}_{S_n}(y)}{\pi(S_n)}.$$

The results in Theorem 2.2.2 generalize to this case as well, whereas those in the next section will replace π_* with $f^{-1}(\mathbb{E}f(\pi(S_0)))$.

2.3 Mixing times

Throughout this section assume that the distance to be studied is of the form

$$\text{dist}(\mathbf{P}^n(x, \cdot), \pi) \leq \hat{\mathbb{E}}_n f(\pi(S_n))$$

for a decreasing function $f : [0, 1] \rightarrow \mathbb{R}_+$. For instance, the total variation, L^p and relative entropy bounds in Theorem 2.2.2 are all of this form. Let $\tau(\epsilon)$ denote the mixing time in this distance, that is, the minimum number of steps to guarantee that this distance is at most ϵ .

Mixing time will be bounded using the f -congestion.

Definition 2.3.1. Given a finite Markov chain, and function $f : [0, 1] \rightarrow \mathbb{R}_+$ non-zero except possibly at 0 and 1, then the f -congestion \mathcal{C}_f and f -congestion profile $\mathcal{C}_f(r)$ are given by

$$\forall A \subset V : \mathcal{C}_f(A) = \frac{\int_0^1 f(\pi(A_u)) du}{f(\pi(A))}, \quad \forall r > 0 : \mathcal{C}_f(r) = \max_{\substack{\pi(A) \leq r, \\ A \neq \emptyset, V}} \mathcal{C}_f(A), \quad \mathcal{C}_f = \mathcal{C}_f(1).$$

The starting point for our calculations will be the following discrete analog of differentiation.

Lemma 2.3.2.

$$\begin{aligned} \hat{\mathbb{E}}_{n+1} f(\pi(S_{n+1})) - \hat{\mathbb{E}}_n f(\pi(S_n)) &= -\hat{\mathbb{E}}_n f(\pi(S_n)) (1 - \mathcal{C}_{af(a)}(S_n)) \\ &\leq -(1 - \mathcal{C}_{af(a)}) \hat{\mathbb{E}}_n f(\pi(S_n)) \end{aligned}$$

Proof. The inequality is because $\forall S \subset V : 1 - \mathcal{C}_{af(a)} \leq 1 - \mathcal{C}_{af(a)}(S)$, by definition of $\mathcal{C}_{af(a)}$. For the equality,

$$\begin{aligned} \hat{\mathbb{E}}_{n+1} f(\pi(S_{n+1})) &= \hat{\mathbb{E}}_n \sum_S \hat{\mathbf{K}}(S_n, S) f(\pi(S)) \\ &= \hat{\mathbb{E}}_n f(\pi(S_n)) \frac{\sum_S \mathbf{K}(S_n, S) \pi(S) f(\pi(S))}{\pi(S_n) f(\pi(S_n))} = \hat{\mathbb{E}}_n f(\pi(S_n)) \mathcal{C}_{af(a)}(S_n) \end{aligned}$$

□

A basic mixing time bound follows easily:

Corollary 2.3.3. *In discrete time*

$$\tau(\epsilon) \leq \left\lceil \frac{1}{1 - \mathcal{C}_{af(a)}} \log \frac{f(\pi_*)}{\epsilon} \right\rceil$$

Proof. By Lemma 2.3.2 $\hat{\mathbb{E}}_{n+1}f(\pi(S_{n+1})) \leq \mathcal{C}_{af(a)} \hat{\mathbb{E}}_n f(\pi(S_n))$. Applying induction to this yields the relation $\hat{\mathbb{E}}_n f(\pi(S_n)) \leq \mathcal{C}_{af(a)}^n f(\pi(S_0))$. Solving for when this drops to ϵ and using the approximation $\log \mathcal{C}_{af(a)} \leq -(1 - \mathcal{C}_{af(a)})$, gives the corollary. \square

This can be generalized to take into consideration set sizes. A stronger bound holds under a fairly weak convexity condition, with about a factor of two lost in the general case.

Theorem 2.3.4. *If $x(1 - \mathcal{C}_{af(a)}(f^{-1}(x)))$ is convex then*

$$\tau(\epsilon) \leq \left[\int_{\pi_*}^{f^{-1}(\epsilon)} \frac{-f'(x) dx}{f(x)(1 - \mathcal{C}_{af(a)}(x))} \right],$$

while in general

$$\tau(\epsilon) \leq \left[\int_{f^{-1}(f(\pi_*)/2)}^{f^{-1}(\epsilon/2)} \frac{-2f'(x) dx}{f(x)(1 - \mathcal{C}_{af(a)}(x))} \right].$$

Proof. First consider the convex case.

By Lemma 2.3.2 and Jensen's inequality for the convex function $x(1 - \mathcal{C}_{af(a)}(f^{-1}(x)))$,

$$\begin{aligned} \hat{\mathbb{E}}_{n+1}f(\pi(S_{n+1})) - \hat{\mathbb{E}}_n f(\pi(S_n)) &= -\hat{\mathbb{E}}_n f(\pi(S_n)) (1 - \mathcal{C}_{af(a)}(S_n)) \\ &\leq -\hat{\mathbb{E}}_n f(\pi(S_n)) [1 - \mathcal{C}_{af(a)}(f^{-1} \circ f(\pi(S_n)))] \\ &\leq -[\hat{\mathbb{E}}_n f(\pi(S_n))] [1 - \mathcal{C}_{af(a)}(f^{-1}(\hat{\mathbb{E}}_n f(\pi(S_n))))]. \end{aligned} \quad (2.1)$$

Since $I(n) = \hat{\mathbb{E}}_n f(\pi(S_n))$ and $1 - \mathcal{C}_{af(a)}(f^{-1}(x))$ are non-increasing, the piecewise linear extension of $I(n)$ to $t \in \mathbb{R}_+$ satisfies

$$I'(t) \leq -I(t) [1 - \mathcal{C}_{af(a)}(f^{-1}(I(t)))] \quad (2.2)$$

At integer t the derivative can be taken from either right or left.

Then,

$$\int_{I(0)}^{I(t)} \frac{dI}{I(1 - \mathcal{C}_{af(a)}(f^{-1}(I)))} \leq - \int_0^t dt = -t.$$

A change of variables to $v = f^{-1}(I)$ implies that

$$\int_{f^{-1}(I(0))}^{f^{-1}(I(t))} \frac{f'(v) dv}{f(v)(1 - \mathcal{C}_{af(a)}(v))} \leq -t.$$

By continuity of $I(t)$ there exists T such that $I(T) = \epsilon$. The theorem follows from $f^{-1}(I(0)) = f^{-1}(f(\pi_*)) = \pi_*$ and $f^{-1}(I(T)) = f^{-1}(\epsilon)$.

For the general case, use Lemma 2.3.5 instead of convexity at (2.1). \square

Lemma 2.3.5. *If $Z \geq 0$ is a nonnegative random variable and g is a nonnegative increasing function, then*

$$E(Zg(Z)) \geq \frac{EZ}{2} g(EZ/2).$$

Proof. See [26]. Let A be the event $\{Z \geq EZ/2\}$. Then $E(Z1_{A^c}) \leq EZ/2$, so $E(Z1_A) \geq EZ/2$. Therefore,

$$E(Zg(2Z)) \geq E(Z1_A g(EZ)) \geq \frac{EZ}{2} g(EZ).$$

Let $U = 2Z$ to get the result. \square

It is fairly easy to translate these to mixing time bounds. For instance, if $f(a) = \sqrt{\frac{1-a}{a}}$ then by Theorem 2.2.2, Corollary 2.3.3 and Theorem 2.3.4 the L^2 -mixing times (denoted by $\tau_2(\epsilon)$) are:

$$\tau_2(\epsilon) \leq \begin{cases} \left\lceil \frac{1}{2(1 - \mathcal{C}_{\sqrt{a(1-a)}})} \log \frac{1 - \pi_*}{\pi_* \epsilon^2} \right\rceil & \text{in general} \\ \left\lceil \int_{\pi_*}^{\frac{1}{1+\epsilon^2}} \frac{dr}{2r(1-r)(1 - \mathcal{C}_{\sqrt{a(1-a)}}(r))} \right\rceil & \text{if } r \left(1 - \mathcal{C}_{\sqrt{a(1-a)}}\left(\frac{1}{1+r^2}\right)\right) \text{ is convex} \\ \left\lceil \int_{\frac{4\pi_*}{1+3\pi_*}}^{\frac{1}{1+\epsilon^2/4}} \frac{dr}{r(1-r)(1 - \mathcal{C}_{\sqrt{a(1-a)}}(r))} \right\rceil & \text{in general} \end{cases}$$

By making the change of variables $x = \frac{r}{1+r}$ and applying a few pessimistic approximations one obtains a result more strongly resembling average conductance bounds:

$$\tau_2(\epsilon) \leq \begin{cases} \left\lceil \frac{1}{1 - \mathcal{C}_{\sqrt{a(1-a)}}} \log \frac{1}{\epsilon \sqrt{\pi_*}} \right\rceil & \text{in general} \\ \left\lceil \int_{\pi_*}^{1/\epsilon^2} \frac{dx}{2x(1 - \mathcal{C}_{\sqrt{a(1-a)}}(x))} \right\rceil & \text{if } x \left(1 - \mathcal{C}_{\sqrt{a(1-a)}}\left(\frac{1}{1+x^2}\right)\right) \text{ is convex} \\ \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{dx}{x(1 - \mathcal{C}_{\sqrt{a(1-a)}}(x))} \right\rceil & \text{in general} \end{cases}$$

It is often unnecessary to compute $\mathcal{C}_f(r)$ for $r > 1/2$. Observe that u almost everywhere $(A^c)_u = (A_{1-u})^c$. It follows that if $f(a) = f(1-a)$ then

$$\mathcal{C}_f(A) = \int_0^1 \frac{f(\pi(A_{1-u}))}{f(\pi(A))} du = \int_0^1 \frac{f(\pi((A_{1-u})^c))}{f(\pi(A^c))} du = \int_0^1 \frac{f(\pi((A^c)_u))}{f(\pi(A^c))} du = \mathcal{C}_f(A^c). \quad (2.3)$$

In particular, $\forall r \geq 1/2 : \mathcal{C}_f(r) = \mathcal{C}_f(1/2) = \max_{\pi(A) \leq 1/2} \mathcal{C}_f(A)$.

Remark 2.3.6. *Mixing time bounds implied by the theorems of this section follow easily for the other distances, but for instance with $\mathcal{C}_{a(1-a)}$ for total variation distance and $\mathcal{C}_{a \log(1/a)}$ for relative entropy. However, it is often better to work with a harder distance, such as bounding total variation mixing ($\tau_{TV}(\epsilon)$) by instead bounding L^2 mixing ($\tau_2(\epsilon)$) and applying the relation $\tau_{TV}(\epsilon) \leq \tau_2(2\epsilon)$. The quantities are related by $\mathcal{C}_{a \log(1/a)}(A) \leq \frac{1}{2}(1 + \mathcal{C}_{a(1-a)}(A))$ (see remarks after Theorem 4.6.2) and $\mathcal{C}_{\sqrt{a(1-a)}}(A) \leq \sqrt{\mathcal{C}_{a(1-a)}(A)}$ (Cauchy-Schwartz), so generally the relative entropy or L^2 -mixing bounds are less than a factor two worse than the total variation bound. In contrast, the lazy walk on a binary cube $\{0, 1\}^d$ has tiny $1 - \mathcal{C}_{a(1-a)}(\{x\}) \approx \frac{d+1}{2} 2^{-d}$, but huge $1 - \mathcal{C}_{\sqrt{a(1-a)}}(\{x\}) \approx \frac{1}{2} - \frac{1}{2\sqrt{d}}$, so the L^2 bounds will give much better asymptotics for this example.*

2.4 Continuous Time

Not much need be changed for continuous time. Let $H_t = e^{-t(1-P)}$ denote the continuous time Markov chain at time t . It is easily verified that if $\hat{K}_t = e^{-t(1-\hat{K})}$ then

$$H_t(x, y) = \hat{\mathbb{E}}_t \pi_{S_t}$$

where $S_0 = \{x\}$ and $\hat{\mathbb{E}}_t$ is the expectation under the walk \hat{K}_t . Bounds involving $P^n(x, y)$ then translate directly into bounds in terms of $H_t(x, y)$. Once Lemma 2.3.2 is replaced by

$$\frac{d}{dt} \hat{\mathbb{E}}_t f(\pi(S_t)) = -\hat{\mathbb{E}}_t f(\pi(S_t))(1 - \mathcal{C}_{af(a)}(S_t))$$

then mixing time bounds also carry over to the continuous-time case, although it is no longer necessary to approximate by a derivative at (2.2) nor necessary to take the ceiling of the bounds.

Chapter 3

Conductance and Modified Conductance

The most common geometric tool for studying mixing time is the *conductance* Φ , a measure of the chance of leaving a set after a single step. Such bounds have been shown only for L^2 mixing time. In this section we show bounds on f -congestion in terms of conductance for lazy walks, the most common situation. The real innovation of this section, however, is the modified conductance, a new quantity which is equivalent to conductance for a lazy walk in L^2 distance, but which also applies to walks with no holding probability and to other distances as well. We finish the section with a discussion on how our evolving set bounds are effected by changes in edge or vertex-expansion, or through re-scaling the transition kernel when slowing down a walk to increase its laziness.

3.1 Conductance

Let us begin with a formal definition of conductance.

Definition 3.1.1. The conductance $\tilde{\Phi}$ and conductance profile $\tilde{\Phi}(r)$ are given by

$$\forall A \subset V : \tilde{\Phi}(A) = \frac{Q(A, A^c)}{\pi(A)\pi(A^c)}, \quad \forall r > 0 : \tilde{\Phi}(r) = \min_{\pi(A) \leq r} \tilde{\Phi}(A), \quad \tilde{\Phi} = \tilde{\Phi}(1/2) = \min_{A \subset V} \tilde{\Phi}(A).$$

The conductance Φ and conductance profile $\Phi(r)$ are defined similarly, but in terms of $\Phi(A) = \frac{Q(A, A^c)}{\min\{\pi(A), \pi(A^c)\}}$. When necessary, notation such as Φ_K will be used to denote conductance for Markov chain K .

The conductance profile $\tilde{\Phi}(r)$ can also be used to upper bound the various f -congestion quantities \mathcal{C}_f when the Markov chain is lazy. The argument is not hard (see also [26]).

Theorem 3.1.2. *Given a lazy Markov chain, and f concave, then*

$$\mathcal{C}_f(A) \leq \frac{f(\pi(A) + 2Q(A, A^c)) + f(\pi(A) - 2Q(A, A^c))}{2f(\pi(A))}.$$

Proof. For a lazy chain, if $u > 1/2$ then $A_u \subset A$, and so

$$\int_{1/2}^1 \pi(A_u) du = \sum_{y \in A} \left(\frac{Q(A, y)}{\pi(y)} - \frac{1}{2} \right) \pi(y) = Q(A, A) - \frac{\pi(A)}{2} = \frac{\pi(A)}{2} - Q(A, A^c).$$

By the Martingale property $\int_0^1 \pi(A_u) du = \pi(A)$ it follows that

$$\int_0^{1/2} \pi(A_u) du = \pi(A) - \int_{1/2}^1 \pi(A_u) du = \frac{\pi(A)}{2} + Q(A, A^c).$$

Recall Jensen's inequality, that $\int g \circ h(u) du \leq g(\int h(u) du)$ if u is a probability distribution and g is concave. By concavity of f ,

$$C_f(A) = \frac{\int_0^{1/2} f(\pi(A_u)) \frac{du}{1/2} + \int_{1/2}^1 f(\pi(A_u)) \frac{du}{1/2}}{2f(\pi(A))} \leq \frac{f\left(\int_0^{1/2} \pi(A_u) \frac{du}{1/2}\right) + f\left(\int_{1/2}^1 \pi(A_u) \frac{du}{1/2}\right)}{2f(\pi(A))}.$$

□

For each choice of f a bit of simplification leads to bounds on C_f . For instance, a lazy Markov chain will have

$$C_{\sqrt{a(1-a)}}(A) \leq \sqrt{1 - \tilde{\Phi}(A)^2} \quad \text{and} \quad \tau_2(\epsilon) \leq \left\lceil \frac{2}{\tilde{\Phi}^2} \log \frac{1}{\epsilon \sqrt{\pi_*}} \right\rceil. \quad (3.1)$$

See proof of Theorem 3.4.1 for a similar calculation. A conductance bound for a non-lazy walk will be considered later.

3.2 Modified conductance

While the conductance has proven useful for studying lazy walks, if the chain is not lazy then the conductance $\tilde{\Phi}(r)$ is not useful for studying mixing. Consider the simple random walk on the complete bipartite graph $K_{m,m}$, a periodic Markov chain. Every subset $A \subset K_{m,m}$ has many edges to A^c so conductance is large, but if A is one of the bipartitions then a Markov chain starting in A will bounce from A to A^c and back again, but it will never mix.

The problem here is that the Markov chain never grows into a larger set, but is always stuck in half of the space. Therefore, it seems more appropriate to consider how much flow from A reaches a strictly larger set, that is the worst flow into a set B where $\pi(B) = \pi(A^c)$. In particular, we consider $\Psi(A) = \Psi(A, \pi(A^c))$ where

$$\Psi(A, t) = \min_{\substack{B \subset V, v \in V, \\ \pi(B) \leq t, \pi(B \cup v) > t}} Q(A, B) + (t - \pi(B)) \frac{Q(A, v)}{\pi(v)} \quad (3.2)$$

is the smallest flow from A to a set of size t . For a lazy chain the minimum in $\Psi(A)$ occurs at $B = A^c$, so $\Psi(A) = Q(A, A^c)$. In general, if π is uniform then $\Psi(A)$ simplifies to $\Psi(A) = \min_{\pi(B)=\pi(A^c)} Q(A, B)$.

It is now possible to define the set quantity that is the main innovation of this section.

Definition 3.2.1. The *modified conductance* $\tilde{\phi}$ and modified conductance profile $\tilde{\phi}(r)$ are given by

$$\tilde{\phi}(A) = \frac{\Psi(A)}{\pi(A)\pi(A^c)}, \quad \tilde{\phi}(r) = \min_{\pi(A) \leq r} \tilde{\phi}(A), \quad \tilde{\phi} = \tilde{\phi}(1/2) = \min_{A \subset V} \tilde{\phi}(A).$$

Define $\phi(A)$ similarly but without $\pi(A^c)$ in the denominator.

For a lazy chain $\Psi(A) = \mathbf{Q}(A, A^c)$ and so $\tilde{\phi}(A) = \tilde{\Phi}(A)$, and modified conductance extends conductance to the non-lazy case. The modified conductance captures important properties quite well. For instance, a connected reversible chain has $\Psi(A) = 0$ if and only if A is one of the bipartitions of a periodic walk; the minimum in $\Psi(A)$ is then achieved by $B = A$, and $\Psi(A) = \mathbf{Q}(A, A) = 0$ rather than $\mathbf{Q}(A, A^c) > 0$ as with conductance.

An alternate interpretation of $\Psi(A)$ is as follows. Given a set $A \subset V$ let $\wp_A \in [0, 1]$ satisfy

$$\inf\{y : \pi(A_y) \leq \pi(A)\} \leq \wp_A \leq \sup\{y : \pi(A_y) \geq \pi(A)\}.$$

The set $V \setminus A_{\wp_A}$ contains the vertices with minimum flow from A , and so if $u < \wp_A$ then $\pi(A_u) - \pi(A) = \pi(\{y \in V \setminus A_{\wp_A} : \mathbf{Q}(A, y) \geq u\pi(y)\})$. It follows that

$$\Psi(A) = \int_0^{\wp_A} (\pi(A_u) - \pi(A)) du = \int_{\wp_A}^1 (\pi(A) - \pi(A_u)) du = \frac{1}{2} \int_0^1 |\pi(A) - \pi(A_u)| du, \quad (3.3)$$

where the first equality is from the definition of $\Psi(A)$ and the second is from Lemma 2.1.3. Since u -almost everywhere $A_u = ((A^c)_{1-u})^c$ the final equality shows that $\Psi(A) = \Psi(A^c)$, a property which is also satisfied by conventional set expansion with $\mathbf{Q}(A, A^c) = \mathbf{Q}(A^c, A)$.

3.3 An Inequality Prover

With this formulation of $\Psi(A)$ it is possible to upper and lower bound each $\mathcal{C}_f(A)$ via Jensen's inequality, although the upper bounds require a careful setup. However, an argument based on Lemma 3.3.1 is more appealing because it gives a general result for all concave f , and it immediately implies sharpness due to the explicit constructions (3.6) and (3.7). We find it also to be more intuitive than Jensen, as it relates a graphical representation of $\pi(A_u)$ directly to the f -congestion.

Lemma 3.3.1. *Suppose that $f : [0, 1] \rightarrow \mathbb{R}$ is concave, and $g, \hat{g} : [0, 1] \rightarrow [0, 1]$ are functions such that \hat{g} is non-increasing and $(g - \hat{g})$ is continuous on a measure one open subset of $(0, 1)$ (e.g. a step function). If*

$$\forall t \in [0, 1] : \int_0^t g(u) du \geq \int_0^t \hat{g}(u) du$$

and $\int_0^1 g(u) du = \int_0^1 \hat{g}(u) du$, then

$$\int_0^1 f \circ g(u) du \leq \int_0^1 f \circ \hat{g}(u) du.$$

Proof. First we use concavity. Given $\delta, c > 0$ and $x, y \in [0, 1]$ with $1 - \delta \geq x \geq y \geq \delta c^{-1}$, let $\lambda = \frac{\delta/c}{x - y + (1 + c^{-1})\delta} \in (0, 1)$. Then $x = \lambda c(y - \delta c^{-1}) + (1 - \lambda)(x + \delta)$ and $y = (1 - \lambda)(y - \delta c^{-1}) + \lambda(x + \delta)$.

By concavity, $f(x) \geq \lambda c f(y - \delta c^{-1}) + (1 - \lambda c) f(x + \delta)$ and $f(y) \geq (1 - \lambda) f(y - \delta c^{-1}) + \lambda f(x + \delta)$. It follows that

$$f(x) + c f(y) \geq f(x + \delta) + c f(y - \delta c^{-1}). \quad (3.4)$$

Now, let $S_1 = \{u \in (0, 1) : g(u) > \hat{g}(u)\}$, $S_2 = \{u \in (0, 1) : g(u) = \hat{g}(u)\}$ and $S_3 = \{u \in (0, 1) : g(u) < \hat{g}(u)\}$. If $(g - \hat{g})$ is continuous on an open interval I then $\{u \in I : g(u) > \hat{g}(u)\}$ and $\{u \in I : g(u) < \hat{g}(u)\}$ are open subsets of I . Hence S_1 and S_3 are open subsets of $(0, 1)$, since $(g - \hat{g})$ is continuous on a measure one countable union of disjoint open intervals. More specifically, each is a countable union of disjoint open intervals. Define

$$h(u) = \sup \left\{ s \in [0, 1] : \int_0^u (g - \hat{g})^+(x) dx \geq \int_0^s (g - \hat{g})^-(x) dx \right\}$$

where $F^\pm(x) = \max\{0, \pm F(x)\}$. Then h is a bijection from $S_1 \setminus S$ to S_3 , where $S = \{u \in S_1 : h(u) \in \overline{S_3} \setminus S_3\}$ is countable. Moreover, when $u \in S_1 \setminus S$ then $h'(u) = \frac{(g - \hat{g})(u)}{(\hat{g} - g)(h(u))} > 0$. Hence, via integration by substitution,

$$\int_{S_3} (f \circ \hat{g})(x) dx = \int_{S_1 \setminus S} (f \circ \hat{g})(h(u)) h'(u) du,$$

and likewise with g in place of \hat{g} . The lemma then follows from integration:

$$\begin{aligned} \int_0^1 f \circ \hat{g}(u) du &= \int_{S_1} + \int_{S_2} + \int_{S_3} (f \circ \hat{g})(u) du \\ &= \int_{S_1 \setminus S} ((f \circ \hat{g})(u) + h'(u) (f \circ \hat{g})(h(u))) du + \int_{S_2} (f \circ \hat{g})(u) du \\ &\geq \int_{S_1 \setminus S} ((f \circ g)(u) + h'(u) (f \circ g)(h(u))) du + \int_{S_2} (f \circ g)(u) du \\ &= \int_{S_1} + \int_{S_2} + \int_{S_3} (f \circ g)(u) du = \int_0^1 f \circ g(u) du \end{aligned}$$

The inequality was an application of (3.4) with $x = \hat{g}(u)$, $y = \hat{g}(h(u))$, $c = h'(u)$, and $\delta = g(u) - \hat{g}(u) = c(\hat{g}(h(u)) - g(h(u)))$. \square

The lemma implies that for any set $A \subset V$, and for some initial conditions, if there are non-increasing step functions $m, M : [0, 1] \mapsto [0, 1]$ such that

$$\begin{aligned} \forall t \in [0, 1] : \int_0^t M(u) du &\geq \int_0^t \pi(A_u) du \geq \int_0^t m(u) du \\ \text{and } \int_0^1 M(u) du &= \int_0^1 \pi(A_u) du = \int_0^1 m(u) du \end{aligned} \quad (3.5)$$

then for every concave function $f(x)$ it follows that

$$\frac{\int_0^1 f(M(u)) du}{f(\pi(A))} \leq \mathcal{C}_f(A) \leq \frac{\int_0^1 f(m(u)) du}{f(\pi(A))}.$$

In the problem at hand, $\pi(A_u) \in [0, 1]$ is non-increasing and equation (3.3) implies $\Psi(A)$ is the area below $\pi(A_u)$ and above $\pi(A)$, and also above $\pi(A_u)$ and below $\pi(A)$. The extreme cases of $\pi(A_u)$ can be drawn immediately, as in Figure 3.1.




Figure 3.1: Distributions such that $\int_0^t M(u) du \geq \int_0^t \pi(A_u) du \geq \int_0^t m(u) du$ given $\Psi(A)$ and φ_A .

3.4 Bounds on f -congestion $\mathcal{C}_f(A)$

We now show modified conductance bounds on some of the f -congestion quantities of interest.

Theorem 3.4.1. *Given a subset $A \subset V$ then*

$$\begin{aligned} \tilde{\phi}(A) &\geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 1 - \sqrt{1 - \tilde{\phi}(A)^2} \geq \tilde{\phi}(A)^2/2 \\ \tilde{\phi}(A) &\geq 1 - \mathcal{C}_{a \log(1/a)}(A) \geq \frac{2\phi(A)^2}{\log(1/\pi(A))} \\ \tilde{\phi}(A) &\geq 1 - \mathcal{C}_{a(1-a)}(A) \geq 4\tilde{\phi}(A)^2\pi(A)(1 - \pi(A)) \end{aligned}$$

Proof. For the upper bound, Figure 3.1 shows that, given $\Psi(A)$ then $\forall t \in [0, 1] : \int_0^t M(u) du \geq \int_0^t \pi(A_u) du$ and $\int_0^1 M(u) du = \pi(A) = \int_0^1 \pi(A_u) du$, where

$$M(u) = \begin{cases} 0 & \text{if } u > 1 - \frac{\Psi(A)}{\pi(A)} \\ \pi(A) & \text{if } u \in \left(\frac{\Psi(A)}{1-\pi(A)}, 1 - \frac{\Psi(A)}{\pi(A)} \right] \\ 1 & \text{if } u \leq \frac{\Psi(A)}{1-\pi(A)} \end{cases} \quad (3.6)$$

By Lemma 3.3.1 any choice of $f(z)$ which is concave and non-negative will therefore satisfy

$$\begin{aligned} \mathcal{C}_f(A) &\geq \frac{\int_0^1 f(M(u)) du}{f(\pi(A))} \\ &= \frac{\Psi(A)}{\pi(A)} \frac{f(0)}{f(\pi(A))} + \left(1 - \frac{\Psi(A)}{\pi(A)\pi(A^c)} \right) \frac{f(\pi(A))}{f(\pi(A))} + \frac{\Psi(A)}{1-\pi(A)} \frac{f(1)}{f(\pi(A))} \\ &\geq 1 - \tilde{\phi}(A) \end{aligned}$$

This shows all of the upper bounds.

To prove lower bounds, suppose φ_A and $\Psi(A)$ are known. Then Figure 3.1 demonstrates that $\forall t \in [0, 1] : \int_0^t \pi(A_u) du \geq \int_0^t m(u) du$ and $\int_0^1 m(u) du = \pi(A) = \int_0^1 \pi(A_u) du$, where

$$m(u) = \begin{cases} \pi(A) - \frac{\Psi(A)}{1-\varphi_A} & \text{if } u > \varphi_A \\ \pi(A) + \frac{\Psi(A)}{\varphi_A} & \text{if } u < \varphi_A \end{cases} \quad (3.7)$$

All that remains is to substitute this into the formula for $1 - \mathcal{C}_f(A)$ for the various $f(x)$ of interest, and then minimize over all possible $\wp_A \in [0, 1]$.

The bound on $1 - \mathcal{C}_{a(1-a)}$ is the easiest. Apply Lemma 3.3.1 with $f(z) = a(1-a)$ to obtain

$$\begin{aligned} \mathcal{C}_{a(1-a)}(A) &\leq \frac{\int_0^1 f(\pi(A_u)) du}{f(\pi(A))} \\ &= \wp_A \frac{\pi(A) + \frac{\Psi(A)}{\wp_A}}{\pi(A)} \frac{1 - \pi(A) - \frac{\Psi(A)}{\wp_A}}{1 - \pi(A)} + (1 - \wp_A) \frac{\pi(A) - \frac{\Psi(A)}{1-\wp_A}}{\pi(A)} \frac{1 - \pi(A) + \frac{\Psi(A)}{1-\wp_A}}{1 - \pi(A)} \\ &= 1 - \frac{\tilde{\phi}(A)^2 \pi(A) \pi(A^c)}{\wp_A(1 - \wp_A)} \leq 1 - 4\tilde{\phi}(A)^2 \pi(A) \pi(A^c) \end{aligned}$$

For the lower bound on $\mathcal{C}_{a \log(1/a)}$ proceed similarly.

$$\begin{aligned} \mathcal{C}_{a \log(1/a)}(A) &\leq -\wp_A \frac{\pi(A) + \frac{\Psi(A)}{\wp_A}}{\pi(A) \log \frac{1}{\pi(A)}} \log \left(\pi(A) + \frac{\Psi(A)}{\wp_A} \right) - (1 - \wp_A) \frac{\pi(A) - \frac{\Psi(A)}{1-\wp_A}}{\pi(A) \log \frac{1}{\pi(A)}} \log \left(\pi(A) - \frac{\Psi(A)}{1 - \wp_A} \right) \\ &= 1 - \frac{\wp_A + \phi(A)}{\log \frac{1}{\pi(A)}} \log \frac{\wp_A + \phi(A)}{\wp_A} - \frac{1 - \wp_A - \phi(A)}{\log \frac{1}{\pi(A)}} \log \frac{1 - \wp_A - \phi(A)}{1 - \wp_A} \end{aligned}$$

Then $(1 - \mathcal{C}_{a \log(1/a)}(A)) \log(1/\pi(A)) \geq g(\wp_A, \phi(A)) \geq 2\phi(A)^2$ by Lemma 3.4.4.

Now for $\mathcal{C}_{\sqrt{a(1-a)}}(A)$. Applying Lemma 3.3.1 and equation (3.7) as before,

$$\begin{aligned} \mathcal{C}_{\sqrt{a(1-a)}}(A) &\leq \wp_A \sqrt{\left(1 + \frac{\Psi(A)}{\wp_A \pi(A)}\right) \left(1 - \frac{\Psi(A)}{\wp_A \pi(A^c)}\right)} + (1 - \wp_A) \sqrt{\left(1 - \frac{\Psi(A)}{(1-\wp_A) \pi(A)}\right) \left(1 + \frac{\Psi(A)}{(1-\wp_A) \pi(A^c)}\right)} \\ &= \sqrt{\left(\wp_A + \tilde{\phi}(A) \pi(A^c)\right) \left(\wp_A - \tilde{\phi}(A) \pi(A)\right)} + \sqrt{\left(1 - \wp_A - \tilde{\phi}(A) \pi(A^c)\right) \left(1 - \wp_A + \tilde{\phi}(A) \pi(A)\right)} \end{aligned}$$

The bound on $\mathcal{C}_{\sqrt{a(1-a)}}(A)$ follows from Lemma 3.4.3 with $X = \wp_A + \tilde{\phi}(A) \pi(A^c)$ and $Y = \wp_A - \tilde{\phi}(A) \pi(A)$. \square

It follows, for instance, that

$$\tau_2(\epsilon) \leq \left\lceil \frac{2}{\tilde{\phi}^2} \log \frac{1}{\epsilon \sqrt{\pi_*}} \right\rceil \quad \text{and} \quad \tau_2(\epsilon) \leq \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{2 dr}{r \tilde{\phi}(r)^2} \right\rceil. \quad (3.8)$$

Conductance can be used to obtain a crude lower bound on the modified conductance.

Lemma 3.4.2. *For an ergodic Markov chain, if $\forall x \in V : \mathbf{P}(x, x) \geq \gamma \in [0, 1]$ and $A \subset V$ then*

$$\tilde{\Phi}(A) \geq \tilde{\phi}(A) \geq \min \left\{ 1, \frac{\gamma}{1 - \gamma} \right\} \tilde{\Phi}(A).$$

Proof. The upper bound $\tilde{\phi}(A) \leq \tilde{\Phi}(A)$ is trivial because $\Psi(A) \leq \mathbf{Q}(A, A^c)$. The minimum in the lower bound is equal to 1 exactly when $\gamma \geq 1/2$, but in this case $\tilde{\phi}(A) = \tilde{\Phi}(A)$, so this case is also trivial. It remains to consider the lower bound when $\gamma < 1/2$.

In the definition of $\Psi(A)$ there is a set B , and one vertex v for which only a $\frac{\pi(A^c) - \pi(B)}{\pi(v)}$ fraction is counted. Extend the state space V to a space V' by splitting v into two vertices v_1 and v_2 , with v_1 of size $\pi(A^c) - \pi(B)$, v_2 with the remainder, and ergodic flows into v_1 and v_2 determined by their respective sizes. Then let $C = B \cup v_1$ be the set where $\Psi(A) = \mathbf{Q}(A, C)$. It follows that

$$\begin{aligned} \Psi(A) &= \mathbf{Q}(A, C \cap A^c) + \mathbf{Q}(A, C \cap A) \\ &\geq \mathbf{Q}(A, C \cap A^c) + \gamma \pi(C \cap A) \\ &\geq \mathbf{Q}(A, C \cap A^c) + \gamma \frac{\mathbf{Q}(A, A^c \setminus C)}{1 - \gamma} \\ &\geq \frac{\gamma}{1 - \gamma} \mathbf{Q}(A, A^c) \end{aligned}$$

The first inequality uses the fact that $\forall v \in A : \mathbf{Q}(v, v) \geq \gamma \pi(v)$ and so $\mathbf{Q}(A, v) \geq \gamma \pi(v)$. The second inequality is because $\pi(C \cap A) = \pi(A^c \setminus C) \geq \frac{\mathbf{Q}((A^c \setminus C)^c, A^c \setminus C)}{1 - \gamma} \geq \frac{\mathbf{Q}(A, A^c \setminus C)}{1 - \gamma}$. \square

The $\gamma/(1 - \gamma)$ factor is introduced when converting $C \cap A$ into a subset of A^c , in short primarily because $\mathbf{Q}(A, A^c)$ is not the correct quantity to work with for non-lazy chains. This induces a mixing bound in terms of conductance for non-lazy walks, but this will be substantially improved on later.

Although Lemma 3.3.1 was just used to show the bounds of Theorem 3.4.1, the arguments are easily modified to use Jensen's Inequality instead (see [21] for just such an approach). However, the upper bound of $\tilde{\phi}(A)$ is somewhat more subtle, and so we give here a proof with Jensen's inequality as well.

Proof of upper bounds in Theorem 3.4.1 using only Jensen's Inequality. In the definition of $\Psi(A)$ there is a set B , and one vertex v for which only a $\frac{\pi(A) - \pi(B)}{\pi(v)}$ fraction is counted. Extend the state space V to a space V' by splitting v into two vertices v_1 and v_2 , with v_1 of size $\pi(A) - \pi(B)$, v_2 with the remainder and flows adjusted accordingly. Then let $C = B \cup v_1$ be the set where $\Psi(A)$ is achieved in the space V' .

Observe that when $u > \wp_A$ then $\pi(A_u) = \sum_{y \in C} \delta_{\mathbf{Q}(A, y) \geq u \cdot \pi(y)} \pi(A) \frac{\pi(y)}{\pi(A)}$. Since $\pi(C) = \pi(A)$ and $f(x) = \sqrt{x(1 - x)}$ is concave then by Jensen's inequality

$$\begin{aligned} \int_{\wp_A}^1 \sqrt{\frac{\pi(A_u)(1 - \pi(A_u))}{\pi(A)(1 - \pi(A))}} du &\geq \int_{\wp_A}^1 \sum_{y \in C} \frac{\pi(y)}{\pi(A)} \sqrt{\frac{\delta_{\mathbf{Q}(A, y) \geq u \cdot \pi(y)} \pi(A) (1 - \delta_{\mathbf{Q}(A, y) \geq u \cdot \pi(y)} \pi(A))}{\pi(A)(1 - \pi(A))}} du \\ &= \int_{\wp_A}^1 \sum_{y \in C} \frac{\pi(y)}{\pi(A)} \delta_{\mathbf{Q}(A, y) \geq u \cdot \pi(y)} du \\ &= \int_{\wp_A}^1 \frac{\pi(A_u)}{\pi(A)} du = 1 - \int_0^{\wp_A} \frac{\pi(A_u)}{\pi(A)} du, \end{aligned}$$

where the final equality uses the Martingale property $\int_0^1 \pi(A_u) du = \pi(A)$. Similarly, when $u \leq \wp_A$ then $\pi(A_u) = \sum_{y \in V' \setminus C} (\pi(A) + \delta_{\mathbf{Q}(A, y) \geq u \cdot \pi(y)} \pi(A^c)) \frac{\pi(y)}{\pi(A^c)}$, and because $\pi(V' \setminus C) = \pi(A^c)$ then

Jensen's inequality can be applied once more to obtain

$$\int_0^{\wp A} \sqrt{\frac{\pi(A_u)(1-\pi(A_u))}{\pi(A)(1-\pi(A))}} du \geq \int_0^{\wp A} \frac{1-\pi(A_u)}{1-\pi(A)} du.$$

Combining these inequalities and rearranging a bit gives the result

$$\begin{aligned} \mathcal{C}_{\sqrt{a(1-a)}}(A) &\geq 1 - \int_0^{\wp A} \frac{(1-\pi(A))\pi(A_u) - \pi(A)(1-\pi(A_u))}{\pi(A)\pi(A^c)} du \\ &= 1 - \int_0^{\wp A} \frac{\pi(A_u) - \pi(A)}{\pi(A)\pi(A^c)} du = 1 - \tilde{\phi}(A). \end{aligned}$$

The only properties used in this argument were the concavity of $f(x) = \sqrt{x(1-x)}$ and that $f(0) = f(1) = 0$. These also hold for $1 - \mathcal{C}_{a(1-a)}(A)$ (with $f(x) = x(1-x)$) and $1 - \mathcal{C}_{a \log(1/a)}(A)$ (with $f(x) = x \log(1/x)$), giving the upper bounds for these. \square

The following two inequalities were used in the proof of Theorem 3.4.1:

Lemma 3.4.3. *If $X, Y \in [0, 1]$ then*

$$g(X, Y) = \sqrt{XY} + \sqrt{(1-X)(1-Y)} \leq \sqrt{1-(X-Y)^2}.$$

Proof. Observe that

$$g(X, Y)^2 = 1 - (X+Y) + 2XY + \sqrt{[1-(X+Y)+2XY]^2 - [1-2(X+Y)+(X+Y)^2]}.$$

Now, $\sqrt{A^2 - B} \leq A - B$ if $A^2 \geq B$, $A \leq \frac{1+B}{2}$ and $A \geq B$ (square both sides to show this). These conditions are easily verified with $A = 1 - (X+Y) + 2XY$ and $B = 1 - 2(X+Y) + (X+Y)^2$, and so

$$\begin{aligned} g(X, Y)^2 &\leq 2[1 - (X+Y) + 2XY] - [1 - 2(X+Y) + (X+Y)^2] \\ &= 1 + 2XY - X^2 - Y^2 = 1 - (X-Y)^2 \end{aligned}$$

\square

Lemma 3.4.4. *If $x \in (0, 1)$ and $y \in [0, 1-x)$ then*

$$g(x, y) = (x+y) \log \frac{x+y}{x} + (1-x-y) \log \frac{1-x-y}{1-x} \geq 2y^2.$$

Proof. Start by seeing what can be shown by differentiation.

$$\begin{aligned} \frac{dg}{dx} &= \log\left(1 + \frac{y}{x}\right) - \frac{y}{x} - \log\left(1 - \frac{y}{1-x}\right) - \frac{y}{1-x} \\ \frac{d^2g}{dx^2} &= y^2 \frac{(x+y)x^2 + (1-x)^2(1-(x+y))}{x^2(1-x)^2(x+y)(1-(x+y))} \geq 0 \end{aligned}$$

It follows that $g(x, y)$ is convex with respect to x , and since $\left.\frac{dg}{dx}\right|_{x=(1-y)/2} \leq 0$ and $\left.\frac{dg}{dx}\right|_{x=1/2} \geq 0$ then the minimum occurs at some $x \in [(1-y)/2, 1/2]$.

To lower bound the minimum we first lower bound $g(x, y)$. By the inequality $f(z) = z \log z + (1 - z) \log(1 - z) \geq -\log 2 + 2(z - 1/2)^2$ when $z \in [0, 1]$ it follows that

$$g(x, y) = f(x + y) - \log(1 - x) + (x + y) \log \frac{1 - x}{x} \geq h(x, y)$$

where $h(x, y) = -\log 2 + 2(x + y - 1/2)^2 - \log(1 - x) + (x + y) \log \frac{1 - x}{x}$. Now,

$$\begin{aligned} \frac{dh}{dx} &= y \left(4 - \frac{1}{x(1-x)} \right) + 4x + \log \frac{1-x}{x} - 2 \\ \frac{d^2h}{dx^2} &= y \frac{1-2x}{x^2(1-x)^2} + 4 - \frac{1}{x(1-x)} \\ \frac{d^2h}{dx^2} \Big|_{x=(1-cy)/2} &= \frac{4cy^2(4-c) + 4c^4y^4}{(1-c^2y^2)^2} \end{aligned}$$

The second derivative is positive when $c \in [0, 1]$, and so $h(x, y)$ is convex in x when $x \in [(1 - y)/2, 1/2]$. However, $\frac{dh}{dx} \Big|_{x=1/2} = 0$ and so $h(x, y) \geq h(1/2, y) = 2y^2$ when $x \in [(1 - y)/2, 1/2]$.

It follows that $g(x, y) \geq \min_{x \in [(1-y)/2, 1/2]} g(x, y) \geq \min_{x \in [(1-y)/2, 1/2]} h(x, y) \geq h(1/2, y) = 2y^2$. \square

3.5 Flow distribution

Now let us look at how flow distribution affects the \mathcal{C} quantities. To do this we assume that we have two Markov chains which differ only in a single characteristic, related to either the level of ergodic flow or the way in which the ergodic flow is distributed among the vertices. The following applications of Lemma 3.3.1 then give a good intuition into what governs mixing.

Corollary 3.5.1. *Suppose that \mathcal{M} and \mathcal{M}' are finite irreducible Markov chains with the same stationary distribution π , that $A \subset V$ with $\Psi(A) = \mathbf{Q}(A, A^c)$ (e.g. any subset if \mathcal{M} is lazy), and that $f : [0, 1] \rightarrow \mathbb{R}^+$ is a concave function. Then*

$$1 - \mathcal{C}_f(A) \geq 1 - \mathcal{C}'_f(A)$$

if either of the following two conditions hold:

- Edge expansion / flow: \mathcal{M}' has smaller pointwise flow than \mathcal{M} , that is

$$\forall v \in A^c : \mathbf{Q}(A, v) \geq \mathbf{Q}'(A, v) \quad \text{and} \quad \forall v \in A : \mathbf{Q}(A^c, v) \geq \mathbf{Q}'(A^c, v),$$

- Vertex expansion: \mathcal{M}' has less well distributed flow than \mathcal{M} , that is

$$\begin{aligned} \forall u \in [0, 1] : \sum_{v \in A^c} \min\{u\pi(v), \mathbf{Q}(A, v)\} &\geq \sum_{v \in A^c} \min\{u\pi(v), \mathbf{Q}'(A, v)\} \\ \text{and } \forall u \in [0, 1] : \sum_{v \in A} \min\{u\pi(v), \mathbf{Q}(A^c, v)\} &\geq \sum_{v \in A} \min\{u\pi(v), \mathbf{Q}'(A^c, v)\} \end{aligned}$$

and moreover $\Psi'(A) = \mathbf{Q}'(A, A^c)$ (e.g. if \mathcal{M}' is lazy) and $\mathbf{Q}(A, A^c) = \mathbf{Q}'(A, A^c)$ (i.e. equal ergodic flows).

The first relation says that, all other things being equal, if each vertex in A^c gets less ergodic flow from A , and vice-versa, then the Evolving set mixing time bound will be worse. The converse does not hold, as the periodic walk on the two-point space has higher edge expansion than the lazy two-point walk, but does not mix.

In order to understand the second case we need to define exactly what is meant by vertex expansion. One reasonable definition is to say that the flow is well distributed if cutting it off at some threshold does not cut off too much, that is if the threshold is u then $\sum_{v \in A^c} \min\{u\pi(v), Q(A, v)\}$ is about the same size as $Q(A, A^c)$, and likewise with a sum over $v \in A$. The corollary then says that, all other conditions being equal, lower vertex expansion leads to a slower mixing time.

The requirement that $\Psi(A) = Q(A, A^c)$ arises from the following lemma.

Lemma 3.5.2. *Given an irreducible Markov chain and $A \subset V$, then $\Psi(A) = Q(A, A^c)$ if and only if $\bigcup_{u > \wp_A} A_u \subsetneq A \subseteq A_{\wp_A}$ for $\wp_A = \inf\{u : \pi(A_u) < \pi(A)\}$.*

Proof. Observe that $\Psi(A) = Q(A, A^c)$ if and only if the set B where the minimum occurs in the definition of $\Psi(A)$ can be taken as $B = A^c$. This happens if and only if $\forall v \in A^c, v' \in A : Q(A, v)/\pi(v) \leq Q(A, v')/\pi(v')$, which is in turn equivalent to $\pi(A_u) \geq \pi(A)$ if and only if $A \subseteq A_u$. This occurs if and only if $A \subseteq A_{\wp_A}$ (observe that $\pi(A_{\wp_A}) \geq \pi(A)$) and $\forall u > \wp_A : A_u \subsetneq A$. \square

This shows that $\Psi(A) = Q(A, A^c)$ if and only if the A_u split into two types, everything at $u > \wp_A$ is in A and everything dropped at $u < \wp_A$ is in A^c . Most properties of lazy Markov chains will hold for sets $A \subset V$ when $\Psi(A) = Q(A, A^c)$.

Proof of Corollary 3.5.1. Figure 3.2 gives a visual “proof” using Lemma 3.3.1.




Figure 3.2: The $\pi(A_u)$ for \mathcal{M} vs. smaller edge expansion \mathcal{M}' , and vs. higher vertex expansion \mathcal{M}' .

Let us now show rigorously that Lemma 3.3.1 can in fact be applied, as suggested by the pictorial representations.

First, edge-expansion.

By Lemma 3.5.2, if $u \leq \wp_A$ then $A_u = A \cup (\bigcup_{v \in A^c} \delta_{Q(A, v) \geq u\pi(v)})$, while if $u > \wp_A$ then $A_u = \bigcup_{v \in A} \delta_{Q(A, v) \geq u\pi(v)}$. The Markov chain \mathcal{M}' also splits into cases of $u \leq \wp_A$ and $u > \wp_A$, because if $v \in A^c$ then $Q'(A, v) \leq Q(A, v) \leq \wp_A \pi(v)$, while if $v \in A$ then $Q'(A, v) \geq Q(A, v) \geq \wp_A \pi(v)$, so we may assume $w(\pi(A)) = \wp_A$ for \mathcal{M}' as well.

First consider the case that $t \leq \wp_A$. If $v \in A^c$ then $Q(A, v)/\pi(v) \geq Q'(A, v)/\pi(v)$, and so $\pi(A_u) \geq \pi(A'_u)$ if $u \leq \wp_A$, and in particular $\forall t \in [0, \wp_A] : \int_0^t \pi(A_u) du \geq \int_0^t \pi(A'_u) du$.

Now consider the case when $t > \wp_A$. If $v \in A$ then $\mathbf{Q}(A, v)/\pi(v) \leq \mathbf{Q}'(A, v)/\pi(v)$ and so for $u \geq \wp_A$ it follows that $\pi(A_u) \leq \pi(A'_u)$, and therefore $\forall t \in [\wp_A, 1] : \int_t^1 \pi(A_u) du \leq \int_t^1 \pi(A'_u) du$. But then, by the Martingale property Lemma 2.1.3,

$$\int_0^t \pi(A_u) du = \pi(A) - \int_t^1 \pi(A_u) du \geq \pi(A) - \int_t^1 \pi(A'_u) du = \int_0^t \pi(A'_u) du$$

We have now established that $\forall t \in [0, 1] : \int_0^t \pi(A_u) du \geq \int_0^t \pi(A'_u) du$, and so the corollary follows from Lemma 3.3.1.

Now, vertex-expansion.

If $t \in [0, 1]$ then

$$\begin{aligned} \int_0^t \pi(A_u) du &= \begin{cases} t - \sum_{v \in A^c} (\mathbf{Q}(A, v) - t\pi(v))^+ & \text{if } t \leq \wp_A \\ \pi(A) - \sum_{v \in A} (\mathbf{Q}(A^c, v) - (1-t)\pi(v))^+ & \text{if } t > \wp_A \end{cases} \\ &= \begin{cases} t - \sum_{v \in A^c} \mathbf{Q}(A, v) - \min\{\mathbf{Q}(A, v), t\pi(v)\} & \text{if } t \leq \wp_A \\ \pi(A) - \sum_{v \in A} \mathbf{Q}(A^c, v) - \min\{\mathbf{Q}(A^c, v), (1-t)\pi(v)\} & \text{if } t > \wp_A \end{cases} \\ &= \begin{cases} t - \mathbf{Q}(A, A^c) + \sum_{v \in A^c} \min\{\mathbf{Q}(A, v), t\pi(v)\} & \text{if } t \leq \wp_A \\ \pi(A) - \mathbf{Q}(A, A^c) + \sum_{v \in A} \min\{\mathbf{Q}(A^c, v), (1-t)\pi(v)\} & \text{if } t > \wp_A \end{cases} \end{aligned}$$

and likewise with $\int_0^t \pi(A'_u) du$. By the conditions on vertex-expansion and the relation $\mathbf{Q}(A, A^c) = \mathbf{Q}'(A, A^c)$ it follows that $\int_0^t \pi(A_u) du \geq \int_0^t \pi(A'_u) du$, and so Lemma 3.3.1 finishes the proof. \square

3.6 The effect of re-scaling on mixing time

One feature of spectral gap/profile approaches to bounding mixing times is that they “scale properly,” in the sense that if the Markov chain is slowed by a factor of two by considering $\frac{1+\mathbf{P}}{2}$ then the mixing time bound also changes by a factor of two, even for non-lazy walks. Conductance bounds don’t immediately have this property, although they can be modified to behave accordingly. The following lemma shows that the f -congestion also “scales properly”, in the sense that increasing the holding probability of a lazy walk also changes the f -congestion in an appropriate way. We also have a bound when the holding probability is decreased below $\gamma = 1/2$, but at this point the walk may start to slow due to periodicity, and so our bound in this case is rather more complicated.

Let us start with the simplest case, re-scaling a walk where $\Psi(A) = \mathbf{Q}(A, A^c)$, such as a lazy walk.

Lemma 3.6.1. *Consider a finite, ergodic Markov chain such that $\forall x \in V : \mathbf{P}(x, x) \geq \gamma \in [0, 1]$. If $A \subset V$ with $\Psi(A) = \mathbf{Q}(A, A^c)$ then*

$$1 - \mathcal{C}_f(A) = 2(1 - \gamma)(1 - \mathcal{C}'_f(A))$$

where \mathcal{C}'_f is the f -congestion for Markov kernel $\mathbf{P}' = \frac{1-\gamma'}{1-\gamma} \mathbf{P} + \left(1 - \frac{1-\gamma'}{1-\gamma}\right) \mathbf{1}$, that is the Markov chain re-scaled to holding probability $1/2$.

The condition that $\Psi(A) = \mathbf{Q}(A, A^c)$ is necessary because, as will be seen in Example 3.7.3, a walk on the complete graph K_m with holding probability γ has

$$1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = \left| \frac{m\gamma - 1}{m - 1} \right|.$$

The lemma will fail for $\gamma < 1/m$, exactly the region for which $\Psi(A) \neq \mathbf{Q}(A, A^c)$.

Proof. Let A'_u denote the evolving sets for \mathbf{P}' and A_u denote those for \mathbf{P} .

◦ Suppose that $\gamma \geq 1/2$. Then

$$A_u = \begin{cases} A'_{u/2(1-\gamma)} & \text{if } u \leq 1 - \gamma \\ A & \text{if } u \in [1 - \gamma, \gamma] \\ A'_{1-(1-u)/2(1-\gamma)} & \text{if } u > \gamma \end{cases}$$

Applying these relations when integrating for $\mathcal{C}_f(A)$ leads to the lemma.

◦ Suppose that $\gamma < 1/2$. Fix any $\wp_A \in [\inf\{u : \pi(A_u) \leq \pi(A)\}, \sup\{u : \pi(A_u) \geq \pi(A)\}]$.

To begin with, if $y \in A$ then because $\mathbf{P}' = \frac{1}{2(1-\gamma)}\mathbf{P} + \left(1 - \frac{1}{2(1-\gamma)}\right)\mathbf{I}$ it follows that

$$\mathbf{Q}'(A, y) = \frac{1}{2(1-\gamma)} \mathbf{Q}(A, y) + \left(1 - \frac{1}{2(1-\gamma)}\right) \pi(y).$$

Therefore, if $y \in A$ then $\mathbf{Q}'(A, y) \geq u\pi(y)$ if and only if

$$\frac{\mathbf{Q}(A, y)}{\pi(y)} \geq \left(u - 1 + \frac{1}{2(1-\gamma)}\right) 2(1-\gamma) = 1 - 2(1-u)(1-\gamma).$$

It follows that $A'_u \cap A = A_{1-2(1-u)(1-\gamma)} \cap A$.

A similar argument holds for the case of $y \in A^c$ and shows that $A'_u \cap A^c = A_{2u(1-\gamma)} \cap A^c$.

Combining the case of $A'_u \cap A^c$ and $A'_u \cap A$ shows that

$$A'_u = (A_{1-2(1-u)(1-\gamma)} \cap A) \cup (A_{2u(1-\gamma)} \cap A^c).$$

If $u > \frac{\wp_A}{2(1-\gamma)}$ then $2u(1-\gamma) > \wp_A$ and by Lemma 3.5.2 it follows that $A'_u \cap A^c = A_{2u(1-\gamma)} \cap A^c = \emptyset$ and so $A'_u = A_{1-2(1-u)(1-\gamma)} \cap A$. If, moreover, $u > 1 - \frac{1-\wp_A}{2(1-\gamma)}$ then $1 - 2(1-u)(1-\gamma) > \wp_A$ and so $A_{1-2(1-u)(1-\gamma)} \subseteq A$, and $A'_u = A_{1-2(1-u)(1-\gamma)}$. However, if $\frac{\wp_A}{2(1-\gamma)} < u \leq 1 - \frac{1-\wp_A}{2(1-\gamma)}$ then $1 - 2(1-u)(1-\gamma) \leq \wp_A$, and so by Lemma 3.5.2 $A'_u \cap A = A_{1-2(1-u)(1-\gamma)} \cap A = A$ and $A'_u = A$. Finally, if $u \leq \frac{\wp_A}{2(1-\gamma)}$ then $1 - 2(1-u)(1-\gamma) \leq \wp_A$ since $\gamma \leq 1/2$, so $A'_u \cap A = A$, and since $2u(1-\gamma) \leq \wp_A$ then by Lemma 3.5.2 $A_{2u(1-\gamma)} \supseteq A$, which implies that $A'_u = A_{2u(1-\gamma)}$.

Putting these all together, we have that

$$A'_u = \begin{cases} A_{2u(1-\gamma)} & \text{if } u \leq \frac{\wp_A}{2(1-\gamma)} \\ A & \text{if } u \in \left(\frac{\wp_A}{2(1-\gamma)}, 1 - \frac{1-\wp_A}{2(1-\gamma)}\right] \\ A_{1-2(1-u)(1-\gamma)} & \text{if } u > 1 - \frac{1-\wp_A}{2(1-\gamma)} \end{cases}$$

Applying these relations when integrating for $\mathcal{C}'_f(A)$ leads to the lemma. \square

One consequence of this is a conductance lower bound on congestion which scales properly. See Lemma 4.2.2 for details.

We finish this section by giving a more complicated re-scaling inequality which applies even when $\Psi(A) \neq \mathbb{Q}(A, A^c)$.

Lemma 3.6.2. *Consider a Markov chain \mathbf{P} with holding probability $\gamma \in [0, 1/2]$, and let $f : [0, 1] \rightarrow \mathbb{R}_+$ be a concave function. Let \mathcal{C}'_f be the f -congestion for Markov kernel $\mathbf{P}' = \frac{1-\gamma'}{1-\gamma}\mathbf{P} + \left(1 - \frac{1-\gamma'}{1-\gamma}\right)\mathbf{1}$, that is the Markov chain re-scaled to holding probability $1/2$. Suppose $h : [0, 1] \rightarrow [0, 1]$ is a decreasing function with $\int_0^1 h(w) dw = \pi(A)$. Then*

$$2(1-\gamma)(1-\mathcal{C}'_f(A)) \geq 1-\mathcal{C}_f(A) \geq 2\gamma \left(1 - \frac{\int_0^1 f(h(w)) dw}{f(\pi(A))}\right),$$

if

$$\int_0^{\gamma/2(1-\gamma)} (h(w) - \pi(A)) dw = \int_{1-\gamma/2(1-\gamma)}^1 (\pi(A) - h(w)) dw \quad (3.9)$$

$$\begin{aligned} \int_0^{1/2} \pi(A'_w) dw &= \int_0^{1/2} h(w) dw \\ \forall t \in [0, 1] : \int_0^t \pi(A'_w) dw &\geq \int_0^t h(w) dw \end{aligned} \quad (3.10)$$

If $\int_0^\gamma + \int_{1-\gamma}^1 \pi(A'_w) dw = 2\gamma\pi(A)$ then by setting $h(w) = \pi(A'_w)$ this implies the lower bound $1-\mathcal{C}_f(A) \geq 2\gamma(1-\mathcal{C}'_f(A))$. If this holds in general then the lemma becomes the much more concise $2(1-\gamma)(1-\mathcal{C}'_f(A)) \geq 1-\mathcal{C}_f(A) \geq 2\min\{\gamma, 1-\gamma\}(1-\mathcal{C}'_f(A))$, and there is no need for the function h . In Example 3.7.2 it is shown that the simple random walk on K_m with holding probability $\gamma \geq 1/m$ has $1-\mathcal{C}_{\sqrt{a(1-a)}}(A) = \frac{m}{m-1}(1-\gamma) = 2(1-\gamma)(1-\mathcal{C}'_f(A))$, and so the upper bound is sharp for every $\gamma > 0$. The lower bound is sharp for the two point space because in Example 3.7.1 it was found that $1-\mathcal{C}_{\sqrt{a(1-a)}}(A) = 2\min\{\gamma, 1-\gamma\}$. More generally, it is sharp when A is a bipartition for the simple random walk on the cycle with an even number of vertices and holding probability γ .

Proof. –First, the upper bound.

If $\gamma \geq 1/2$ then by Lemma 3.6.1 the upper bound is an equality, so we may assume $\gamma < 1/2$. By the relations shown in the proof of Lemma 3.6.1, and because we may assume without loss that $\varnothing_A \in [\gamma, 1-\gamma]$, then

$$A_u = \begin{cases} A'_{u/2(1-\gamma)} & \text{if } u \leq \gamma \\ (A'_{u/2(1-\gamma)} \setminus A) \cup A'_{1-(1-u)/2(1-\gamma)} & \text{if } u \in [\gamma, 1-\gamma] \\ A'_{1-(1-u)/2(1-\gamma)} & \text{if } u \geq 1-\gamma \end{cases} \quad (3.11)$$

Integrate over $u \notin [\gamma, 1-\gamma]$:

$$\int_0^\gamma + \int_{1-\gamma}^1 f(\pi(A)) - f(\pi(A_u)) du = 2(1-\gamma) \int_0^{\frac{\gamma}{2(1-\gamma)}} + \int_{1-\frac{\gamma}{2(1-\gamma)}}^1 f(\pi(A)) - f(\pi(A'_w)) dw$$

Now integrate over $u \in [\gamma, 1 - \gamma]$. To do this, recall from equation (3.4) that if $a + b = c + d$ and $a > c > d > b$ then $f(a) + f(b) \leq f(c) + f(d)$. Well, $\pi(A'_{\frac{u}{2(1-\gamma)}}) = \pi(A) + \pi(A'_{\frac{u}{2(1-\gamma)}} \setminus A)$ and so

$$f(\pi(A'_{\frac{u}{2(1-\gamma)}})) + f(\pi(A'_{1-\frac{1-u}{2(1-\gamma)}})) \leq f(\pi(A)) + f(\pi((A'_{\frac{u}{2(1-\gamma)}} \setminus A) \cup A'_{1-\frac{1-u}{2(1-\gamma)}})).$$

It follows that

$$\begin{aligned} & \int_{\gamma}^{1-\gamma} f(\pi(A)) - f(\pi(A_u)) du \\ & \leq \int_{\gamma}^{1-\gamma} \left(f(\pi(A)) - f(\pi(A'_{\frac{u}{2(1-\gamma)}})) \right) + \left(f(\pi(A)) - f(\pi(A'_{1-\frac{1-u}{2(1-\gamma)}})) \right) du \\ & = 2(1-\gamma) \int_{\frac{\gamma}{2(1-\gamma)}}^{1-\frac{\gamma}{2(1-\gamma)}} f(\pi(A)) - f(\pi(A'_w)) dw \end{aligned}$$

Adding the integrals for $u \notin [\gamma, 1 - \gamma]$ and $u \in [\gamma, 1 - \gamma]$ gives the upper bound.

–Now, the lower bound.

The conditions on $h(u)$ were chosen so that if

$$h_2(u) = \begin{cases} h\left(\frac{u}{2(1-\gamma)}\right) & \text{if } u \leq \gamma \\ h\left(\frac{\gamma}{2(1-\gamma)} + \frac{u-\gamma}{2(1-\gamma)}\right) - \pi(A) + h\left(\frac{1}{2} + \frac{u-\gamma}{2(1-\gamma)}\right) & \text{if } u \in [\gamma, 1 - \gamma] \\ h\left(1 - \frac{1-u}{2(1-\gamma)}\right) & \text{if } u > 1 - \gamma \end{cases}$$

then h_2 is decreasing, $\int_0^1 h_2(u) du = \pi(A)$, and $\forall t \in [0, 1] : \int_0^t \pi(A_u) du \geq \int_0^t h_2(u) du$ by (3.10) and (3.11). It follows from Lemma 3.3.1 that

$$\mathcal{C}_f(A) \leq \frac{\int_0^1 f(h_2(u)) du}{f(\pi(A))}.$$

By Jensen's Inequality, the relation $\int_0^1 h(u) du = \pi(A)$ and (3.9),

$$\int_{\gamma}^{1-\gamma} f(h_2(u)) \frac{du}{1-2\gamma} \leq f\left(\int_{\gamma}^{1-\gamma} h_2(u) \frac{du}{1-2\gamma}\right) = f(\pi(A)),$$

and so

$$\begin{aligned} & \int_0^1 f(\pi(A)) - f(h_2(u)) du \\ & \geq \int_0^{\gamma} + \int_{1-\gamma}^1 f(\pi(A)) - f(h_2(u)) du \\ & = 2(1-\gamma) \int_0^{\frac{\gamma}{2(1-\gamma)}} + \int_{1-\frac{\gamma}{2(1-\gamma)}}^1 f(\pi(A)) - f(h(u)) du \\ & = 2\gamma \int_0^1 f(\pi(A)) - f(h_3(u)) du \\ & \geq 2\gamma \int_0^1 f(\pi(A)) - f(h(w)) dw. \end{aligned}$$

In the final equality

$$h_3(u) = \begin{cases} h\left(\frac{\gamma}{1-\gamma}u\right) & \text{when } u \leq 1/2 \\ h\left(1 - \frac{\gamma}{1-\gamma}(1-u)\right) & \text{when } u > 1/2 \end{cases}$$

The final inequality is an application of Lemma 3.3.1. To see this, first apply (3.9) to see that $\int_0^1 h_3(u) du = \pi(A) = \int_0^1 h(u) du$. If $t \leq 1/2$ then $\int_0^t h_3(u) du \geq \int_0^t h(u) du$ because $h(u)$ is a decreasing function and so $h_3(u) \geq h(u)$. If $t > 1/2$ then $\int_0^t h_3(u) du = \pi(A) - \int_t^1 h_3(u) du \geq \pi(A) - \int_t^1 h(u) du = \int_0^t h(u) du$, again because $h(u)$ is decreasing. \square

3.7 Examples

The purpose of this section is to demonstrate sharpness of bounds. We start with the elementary example of a walk on a complete graph, in which each bound is either sharp or at least asymptotically of the correct order. This is followed by a careful analysis of random walk on a cycle, in which we show fairly sharp total variation mixing time bounds. We finish by discussing the simple random walk on a directed non-lazy Eulerian graph, for which our methods appear to give the first proof of a mixing time bound.

First, we see that the conductance bounds are sharp.

Example 3.7.1. Consider the uniform two-point space $\{0, 1\}$ with transition kernel $P(0, 0) = P(1, 1) = \gamma \in [0, 1]$ and $P(0, 1) = P(1, 0) = (1 - \gamma)$. Then $\tilde{\Phi}(A) = 2(1 - \gamma)$, and so by Lemma 4.2.2

$$2(1 - \gamma) \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}} \geq 2 \min\{\gamma, 1 - \gamma\}.$$

Hence $1 - \mathcal{C}_{\sqrt{a(1-a)}} = 2(1 - \gamma)$ if $\gamma \geq 1/2$.

More generally, $1 - \mathcal{C}_{\sqrt{a(1-a)}} \leq \tilde{\phi} = 2 \min\{\gamma, 1 - \gamma\}$ and so the upper and lower bound are equal and $1 - \mathcal{C}_{\sqrt{a(1-a)}} = 2 \min\{\gamma, 1 - \gamma\}$ for all $\gamma \in [0, 1]$.

Theorem 3.4.1 can lead to sharp bounds, even for holding probability under $1/2$.

Example 3.7.2. Consider the random walk on the complete graph K_m with $P(x, y) = 1/m$. Then $\forall A \subset V : \tilde{\phi}(A) = 1$ and so $1 \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 1 - \sqrt{1 - 1^2} = 1$. Moreover, when $\pi(A) = 1/2$ then $1 \geq 1 - \mathcal{C}_{a(1-a)}(A) \geq 1$ and $1 \geq 1 - \mathcal{C}_{a \log(1/a)}(A) \geq (2 \log 2)^{-1} \approx 0.72$. Therefore at least two of the three bounds in Theorem 3.4.1 can be sharp.

By Lemma 3.6.1 a rescaling argument can be used to extend this to sharp bounds for other holding probabilities, as long as $\Psi(A) = Q(A, A^c)$. In particular, if $\gamma \geq 1/m$ then the walk on K_m with $P(x, x) = \gamma$ and $P(x, y) = \frac{1-\gamma}{m-1}, \forall y \neq x$ satisfies $\Psi(A) = \pi(A)\pi(A^c)\frac{m}{m-1}(1-\gamma) = Q(A, A^c)$. Hence, if $\gamma = 1/m$, and P' is the walk with holding probability $1/2$, then $1 - \mathcal{C}'_{\sqrt{a(1-a)}}(A) =$

$$\frac{1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)}{2(1-1/m)} = \frac{1}{2(1-1/m)}.$$

$$1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = 2(1 - \gamma)(1 - \mathcal{C}'_{\sqrt{a(1-a)}}(A)) = \frac{m}{m-1}(1 - \gamma).$$

In fact, the f -congestion can be used to show sharp mixing time bounds, regardless of holding probability.

Example 3.7.3. Given $\alpha \in [-\frac{1}{m-1}, 1]$ consider the walk on K_m with $P(x, y) = (1 - \alpha)/m$ for all $y \neq x$ and $P(x, x) = \alpha + (1 - \alpha)/m$, that is, choose a point uniformly at random and move there with probability $1 - \alpha$, otherwise do nothing.

The n step distribution is $P^n(x, x) = \frac{1}{m} + \alpha^n (1 - \frac{1}{m})$ and $P^n(x, y) = \frac{1}{m} - \frac{\alpha^n}{m}$ for all $y \neq x$. Therefore, when $\alpha \in [0, 1]$ then $D(P^n(x, \cdot) \| \pi) = (1 + o_m(1))\alpha^n \log m$ as $m \rightarrow \infty$. When $\alpha \in [-\frac{1}{m-1}, 1]$ then $\|P^n(x, \cdot) - \pi\|_{TV} = |\alpha|^n (1 - 1/m)$ and $\|P^n(x, \cdot) - \pi\|_{L^2(\pi)} = |\alpha|^n \sqrt{m-1}$.

Now for evolving sets.

If $\alpha \in [0, 1]$ then

$$\pi(A_u) = \begin{cases} 0 & \text{if } u \in (\alpha + (1 - \alpha)\pi(A), 1] \\ \pi(A) & \text{if } u \in ((1 - \alpha)\pi(A), \alpha + (1 - \alpha)\pi(A)] \\ 1 & \text{if } u \in [0, (1 - \alpha)\pi(A)] \end{cases}$$

A quick calculation shows that $\mathcal{C}_{a(1-a)} = \mathcal{C}_{a \log(1/a)} = \mathcal{C}_{\sqrt{a(1-a)}} = \alpha$, and so Theorem 2.3.3 implies $\|P^n(x, \cdot) - \pi\|_{TV} \leq \alpha^n (1 - 1/m)$, $D(P^n(x, \cdot) \| \pi) \leq \alpha^n \log m$ and $\|P^n(x, \cdot) - \pi\|_{L^2(\pi)} \leq \alpha^n \sqrt{m-1}$. Total variation and L^2 bounds are correct, while relative entropy is asymptotically correct.

When $\alpha \in [-\frac{1}{m-1}, 0)$ then

$$\pi(A_u) = \begin{cases} 0 & \text{if } u > (1 - \alpha)\pi(A), \\ \pi(A^c) & \text{if } u > \alpha + (1 - \alpha)\pi(A), \\ 1 & \text{otherwise} \end{cases}$$

This time $\mathcal{C}_{a(1-a)} = \mathcal{C}_{\sqrt{a(1-a)}} = -\alpha$ and so $\|P^n(x, \cdot) - \pi\|_{TV} \leq (-\alpha)^n (1 - 1/m)$ and $\|P^n(x, \cdot) - \pi\|_{L^2(\pi)} \leq (-\alpha)^n \sqrt{m-1}$, both exact.

A harder walk to bound is the simple random walk on the cycle C_m , that is $P(x, x \pm 1) = 1/2$. A bound must distinguish between the (periodic) walk on a cycle of even length, and the (convergent) walk on a cycle of odd length.

Example 3.7.4. The walk on a cycle C_m of even length has $\tilde{\phi} = 0$ because it is bipartite, with the worst set A given by choosing $m/2$ alternating points around the cycle, and $B = A$ in the definition of $\Psi(A)$. Therefore $0 = \tilde{\phi} \geq 1 - \mathcal{C}_f \geq 0$ for all of the quantities dealt with in Theorem 3.4.1. Correctly, none of our bounds show mixing.

Now for the cycle C_m of odd length. If $\pi(A) < 1/2$ then $\Psi(A) \geq 1/2m$, with the worst sets given by points alternating around the cycle, as in the white vertices of Figure 3.3. Then $\Psi(A) = Q(A, B)$ when B contains those points at least distance two from A , one point adjacent to these and A , and the points in A , corresponding to the circled regions in Figure 3.3.

Therefore

$$1 - \mathcal{C}_{a(1-a)}(A) \geq 4\tilde{\phi}^2(A)\pi(A)\pi(A^c) \geq \frac{1}{m^2\pi(A)\pi(A^c)}$$

By Theorem 2.3.4 it follows that if $\epsilon \geq 1/2$ then

$$\tau(\epsilon) \leq \int_{\pi_*}^{1-\epsilon} \frac{dx}{(1-x)(1-\mathcal{C}_{a(1-a)}(x))} \leq \int_{1/m}^{1-\epsilon} m^2 x dx = \frac{m^2}{2} ((1-\epsilon)^2 - 1/m^2)$$




Figure 3.3: Let A be the white vertices and B be the circled points. Then $\Psi(A) = Q(A, B) = 1/2m$.

and so if $x \in V$ then

$$\|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} \leq 1 - \frac{1}{m} \sqrt{1 + 2n} \quad \text{if } n \leq \frac{m^2}{8} - \frac{1}{2}. \quad (3.12)$$

Standard techniques give poor bounds for large epsilon, such as $\epsilon > 1/2$ above.

Bounds for $\epsilon < 1/2$ can be obtained similarly, but better asymptotics can be derived by a slight modification of the argument. Observe that

$$\begin{aligned} \|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} &\leq \frac{1}{\pi(x)} \mathbb{E} \pi(S_n) (1 - \pi(S_n)) \leq \frac{1}{\pi(x)} \frac{\pi_*(1 - \pi_*)}{\sin(3.14 \pi_*)} \mathbb{E} \sin(3.14 \pi(S_n)) \\ &\leq \frac{\sin(3.14 \pi(x))}{\pi(x)} \frac{\pi_*(1 - \pi_*)}{\sin(3.14 \pi_*)} \mathcal{C}_{\sin(\pi a)}^n \leq (1 - \pi_*) \mathcal{C}_{\sin(\pi a)}^n \end{aligned}$$

where 3.14 is used to represent the number π . The choice of $\mathcal{C}_{\sin(\pi a)}$ is because if $\Psi(A) \geq C$ for some constant C then \mathcal{C}_f is minimized by $f(a) = \sin(\pi a)$ (see [21] for details).

Now, when $\wp_A = 1/2$ then by Lemma 3.3.1 and equation (3.7) it follows that $\mathcal{C}_{\sin(\pi a)}(A) \leq \cos(2\pi\Psi(A))$. On the cycle, if $A \subset V$ then $\pi(A_u) > \pi(A)$ when $u < 1/2$, while $\pi(A_u) < \pi(A)$ when $u > 1/2$, so $\wp_A = 1/2$. Combined with the earlier bound $\Psi(A) \geq 1/2m$ it follows that $\mathcal{C}_{\sin(\pi a)}(A) \leq \cos(\pi/m)$. Then

$$\|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} \leq (1 - \pi_*) \mathcal{C}_{\sin(\pi a)}^n = (1 - 1/m) \cos^n(\pi/m). \quad (3.13)$$

A fairly close lower bound holds as well. Let $\lambda_{max} = \max\{\lambda_2, |\lambda_m|\}$ be the second largest magnitude of an eigenvalue of \mathbf{P} . It is easily verified that $\cos\left(\frac{\pi(m-1)}{m}\right)$ is an eigenvalue with eigenvector $f(j) = \cos\left(\frac{2\pi(m-1)j}{m}\right)$, so $\lambda_{max} \geq \left|\cos\left(\frac{\pi(m-1)}{m}\right)\right| = \cos(\pi/m)$. But then

$$\max_{x \in V} \|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} \geq \frac{1}{2} \lambda_{max}^n \geq \frac{1}{2} \cos^n(\pi/m). \quad (3.14)$$

The first inequality is a general bound for time-reversible chains.

One bound that appears in the literature is

$$\frac{1}{2} \cos^n(\pi/m) \leq \max_{x \in V} \|\mathbf{P}^n(x, \cdot) - \pi\|_{TV} \leq e^{-\pi^2 n / 2m^2} \quad \text{if } n \geq m^2/40.$$

Our bound (3.13) is at most $(1 - 1/m)e^{-\pi^2 n/2m^2}$, mildly better overall and with no conditions on n . The old bound also required knowledge of the complete spectrum of the transition matrix. In contrast, we required only examination of edge expansion properties.

We finish with an example where our methods give the only known mixing time bounds, the simple random walk on a directed Eulerian graph.

Example 3.7.5. Consider a directed Eulerian graph with vertex set V and m edges, that is, a strongly connected graph with in-degree=out-degree at each vertex. The simple random walk is a walk which chooses a neighboring vertex uniformly and then transitions there. This walk has $P(x, y) = 1/\text{deg}(x)$ if there is an edge from x to y , and stationary distribution $\pi(x) = \text{deg}(x)/m$. It is known that the lazy simple random walk (i.e. $P(x, x) = 1/2$ and $P(x, y) = 1/2\text{deg}(x)$) has mixing time $\tau_2(\epsilon) = O(m^2 \log(m/\epsilon))$, but nothing seems to be known about the non-lazy simple random walk even on undirected graphs.

Before stating a mixing bound we must exclude graphs on which the simple random walk does not converge. For instance, a bipartite graph. More generally, the walk is non-convergent if a directed graph has k (equal sized) components such that a transition starting in component i always goes to component $i+1 \pmod k$. The problem here is that a walk starting in one component has a neighborhood the same size as the original set, so it never grows to cover the entire space. If we let $N(A) = \{x \in V : Q(A, x) > 0\}$ denote the neighborhood of A , then the following weak expansion condition will suffice to rule out such situations:

$$\forall A \subset V, \pi(A) \leq 1/2, \forall v \in V : \pi(N(A) \setminus v) \geq \pi(A) \quad (3.15)$$

This just says that if any single vertex in the neighborhood of A is removed, then the neighborhood is still at least as big as A . Note this cannot be satisfied if some vertex has only one outgoing edge, and so $\pi_* = \min_{v \in V} \pi(v) \geq 2/m$.

We now lower bound $\Psi(A)$. Suppose $A \subset V$ with $\pi(A) \leq 1/2$, and if $\Psi(A) = Q(A, B) + (\pi(A^c) - \pi(B)) \frac{Q(A, v)}{\pi(v)}$. If $N(A) \subseteq B^c$ then $\pi(N(A) \setminus v) \leq \pi(B^c \setminus v) = 1 - \pi(B \cup v) < \pi(A)$, contradicting the expansion condition. Hence, $N(A) \cap B \neq \emptyset$ and so there are vertices $x \in A, y \in B$ with $P(x, y) > 0$. Then

$$\Psi(A) \geq Q(A, B) \geq \pi(x)P(x, y) = \frac{\text{deg}(x)}{m} \frac{1}{\text{deg}(x)} = \frac{1}{m}.$$

It follows that $\tilde{\phi}(r) \geq \frac{1}{mr(1-r)}$ if $r \leq 1/2$, and so from the convex version of equation (3.8) that

$$\begin{aligned} \tau_2(\epsilon) &\leq \left[\int_{2/m}^{1/2} \frac{dr}{2r(1-r)\tilde{\phi}(r)^2/2} + \int_{1/2}^{\frac{1}{1+\epsilon^2}} \frac{dr}{2r(1-r)\tilde{\phi}^2/2} \right] \\ &\leq \left[\frac{m^2}{12} + \frac{m^2}{8} \log \frac{1}{\epsilon} \right]. \end{aligned}$$

The same argument can be used to improve on the classical $\tau_2(\epsilon) = O(m^2 \log(m/\epsilon))$ bound for the lazy simple walk. Every lazy walk has $\Psi(A) = Q(A, A^c)$, and so $\Psi(A) \geq 1/2m$ even without the expansion condition. It follows that $\tilde{\phi}(r) \geq \frac{1}{2mr(1-r)}$, and so the lazy simple random walk mixes in

$$\tau_2(\epsilon) \leq \left[\frac{m^2}{3} + \frac{m^2}{2} \log \frac{1}{\epsilon} \right].$$

Note that the (lazy or non-lazy) simple random walk on a cycle with an odd number of vertices has $\tau_2(\epsilon) = \Theta(m^2 \log \frac{1}{\epsilon})$, and so even for the lazy simple random walk our bounds are the first ones of the correct order.

A total variation bound can be found by integrating the appropriate total variation result of Theorems 2.3.4 and 3.4.1. Instead, to give a taste of what improvements can be made, we note that in [17] the above technique is sharpened to show that the (non-lazy) simple random walk satisfies

$$\tau_{TV}(\epsilon) \leq \left[\frac{1}{-\log \cos \frac{2\pi}{m}} \log \frac{1-2/m}{\epsilon} \right] \approx \frac{m^2}{2\pi^2} \log \frac{1}{\epsilon}.$$

This bound is exact for the simple random walk on a cycle with 3 vertices (i.e. K_3 with $\alpha = -1/2$ in Example 3.7.3), while more generally equation (3.14) shows an extremely close lower bound for a cycle with an odd number of vertices:

$$\tau_{TV}(\epsilon) \geq \left[\frac{1}{-\log \cos \frac{2\pi}{m}} \log \frac{1}{2\epsilon} \right] \approx \frac{m^2}{2\pi^2} \log \frac{1}{2\epsilon}$$

Numerous other improvements and generalizations are possible. See [17] in which we sharpen this analysis further, extend it to show bounds on other walks such as the max-degree walk, and also give near-optimal bounds for spectral gap and other quantities of interest.

Chapter 4

A comparison to previous isoperimetric bounds

How do our new results compare to previous isoperimetric bounds? In this section we compare our new Evolving set mixing bounds to earlier Evolving Set bounds, to Spectral profile bounds, to Blocking Conductance results, and to Discrete Gradient methods.

4.1 Evolving Sets

Morris and Peres' used a more probabilistic argument than ours to show that if $x \in V$ and $S_0 = \{x\}$ then

$$\|\mathbb{P}^n(x, \cdot) - \pi\|_{L^2(\pi)} \leq \hat{\mathbb{E}}_n \frac{\min\{\sqrt{\pi(S_n)}, \sqrt{1 - \pi(S_n)}\}}{\pi(S_n)},$$

not a major difference but up to $\sqrt{2}$ times weaker than our bound in Theorem 2.2.2. They did not have bounds on total variation or relative entropy.

Our rate of contraction $\mathcal{C}_{\sqrt{a(1-a)}}$ on L^2 distance is also better than the $\mathcal{C}_{\sqrt{a}}$ that they showed. Let $f(x, y) = \sqrt{\frac{x}{y}} - \sqrt{\frac{x(1-x)}{y(1-y)}}$ with domain $x, y \in (0, 1)$. This is convex in x because $\frac{d^2}{dx^2} f(x, y) = \frac{1 - (1-x)^{3/2} \sqrt{1-y}}{4(x(1-x))^{3/2} \sqrt{y(1-y)}} \geq 0$. Then by Jensen's inequality,

$$\int_0^1 \left(\sqrt{\frac{\pi(A_u)}{\pi(A)}} - \sqrt{\frac{\pi(A_u)(1 - \pi(A_u))}{\pi(A)(1 - \pi(A))}} \right) du \geq f \left(\int_0^1 \pi(A_u) du, \pi(A) \right) = f(\pi(A), \pi(A)) = 0,$$

showing that $\mathcal{C}_{\sqrt{a(1-a)}}(A) \leq \mathcal{C}_{\sqrt{a}}(A)$.

4.2 Spectral Profile

Two isoperimetric bounds on mixing time are shown in the Spectral Profile paper [9]:

$$\tau_2(\epsilon) \leq \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{4 dr}{r \Phi_{\text{PP}^*}(r)^2} \right\rceil \quad \text{and} \quad \tau_2(\epsilon) \leq \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{2 dr}{\frac{\gamma}{1-\gamma} r \Phi(r)^2} \right\rceil \quad (4.1)$$

The holding probability $\gamma \in [0, 1]$ is such that $\forall x \in V : P(x, x) \geq \gamma$.

It will be shown below that the Evolving set L^2 bounds are at least as good as the bounds of (4.1). However, Evolving set bounds have the advantage that they apply to other distances, such as total variation and relative entropy, for which the isoperimetric bounds on mixing via Spectral gap/profile are merely induced from the L^2 mixing bound. See [17] for an example where modified conductance is used to show a total variation mixing bound which is strictly better than the L^2 mixing bound. Comparison methods involving conductance are also available with Evolving sets, as will be shown in Section 4.4. On the other hand, we do not know of decomposition results which compare favorably to those available with spectral gap or log-Sobolev methods.

First, we show that bounding mixing time with modified conductance is no worse than using the multiplicative reversibilization PP^* in (4.1), but may give a substantial improvement. In particular, it is fairly simple to construct a distribution for $Q(A, y)$ in the proof of Lemma 4.2.1 for which the upper bound is an equality, and likewise for the the lower bound, and so in the worst case scenario $\tilde{\Phi}_{PP^*}(r)^2 = \tilde{\phi}(r)^4$, and the first bound of (4.1) may be nearly as bad as the square of the modified conductance mixing bound.

Lemma 4.2.1.

$$\sqrt{\tilde{\Phi}_{PP^*}(A)} \geq \tilde{\phi}(A) \geq 1 - \sqrt{1 - \tilde{\Phi}_{PP^*}(A)} \geq \frac{1}{2} \tilde{\Phi}_{PP^*}(A).$$

Proof. To simplify notation, in the definition of $\Psi(A)$ assume that the set B satisfies $\pi(B) = \pi(A^c)$, i.e. $\Psi(A) = Q(A, B)$. For the general case it suffices to split an appropriate vertex v , as in the proof of Lemma 3.4.2.

To begin with, we need a few identities:

$$\begin{aligned} Q(A, B) &= \pi(A) - Q(A, B^c) = \pi(A) - (\pi(B^c) - Q(A^c, B^c)) = Q(A^c, B^c) \\ Q_{PP^*}(A, A^c) &= \sum_{y \in V} Q(A, y) P^*(y, A^c) = \sum_{y \in V} \frac{Q(A, y)}{\pi(y)} \left(1 - \frac{Q(A, y)}{\pi(y)}\right) \pi(y) \end{aligned}$$

First we bound the terms in the summation for $Q_{PP^*}(A, A^c)$ for which $y \in B$.

Let $f(a) = a(1 - a)$. The proof can be completed using Lemma 3.3.1, but using Jensen's Inequality is much simpler.

$$\begin{aligned} \sum_{y \in B} \frac{Q(A, y)}{\pi(y)} \left(1 - \frac{Q(A, y)}{\pi(y)}\right) \frac{\pi(y)}{\pi(B)} &= \sum_{y \in B} f\left(\frac{Q(A, y)}{\pi(y)}\right) \frac{\pi(y)}{\pi(B)} \\ &\leq f\left(\sum_{y \in B} \frac{Q(A, y)}{\pi(y)} \frac{\pi(y)}{\pi(B)}\right) \\ &= f\left(\frac{Q(A, B)}{\pi(B)}\right) = \frac{\Psi(A)}{\pi(A^c)} \left(1 - \frac{\Psi(A)}{\pi(A^c)}\right) \end{aligned}$$

To show a lower bound, note that if $y \in B$ and $v \notin B$ then $\frac{Q(A, y)}{\pi(y)} \leq \frac{Q(A, v)}{\pi(v)}$ by definition of set B , and so $\frac{Q(A, y)}{\pi(y)} \leq \frac{Q(A, B^c)}{\pi(B^c)} = \frac{\pi(A) - \Psi(A)}{\pi(A)} = 1 - \frac{\Psi(A)}{\pi(A)}$. Then

$$\begin{aligned} \sum_{y \in B} \frac{Q(A, y)}{\pi(y)} \left(1 - \frac{Q(A, y)}{\pi(y)}\right) \frac{\pi(y)}{\pi(B)} &\geq \sum_{y \in B} \frac{Q(A, y)}{\pi(y)} \frac{\Psi(A)}{\pi(A)} \frac{\pi(y)}{\pi(B)} \\ &= \frac{\Psi(A)}{\pi(A)} \frac{Q(A, B)}{\pi(B)} = \frac{\Psi(A)^2}{\pi(A)\pi(A^c)} \end{aligned}$$

To bound the terms over B^c use the relation $\frac{Q(A,y)}{\pi(y)} = 1 - \frac{Q(A^c,y)}{\pi(y)}$ to re-write the sum:

$$\sum_{y \in B^c} \frac{Q(A,y)}{\pi(y)} \left(1 - \frac{Q(A,y)}{\pi(y)}\right) \frac{\pi(y)}{\pi(B^c)} = \sum_{y \in B^c} \frac{Q(A^c,y)}{\pi(y)} \left(1 - \frac{Q(A^c,y)}{\pi(y)}\right) \frac{\pi(y)}{\pi(B^c)}$$

Then follow the same steps as before, and apply the relation $Q(A^c, B^c) = \Psi(A)$, to obtain

$$\Psi(A) \left(1 - \frac{\Psi(A)}{\pi(A)}\right) \geq \sum_{y \in B^c} \frac{Q(A^c,y)}{\pi(y)} \left(1 - \frac{Q(A^c,y)}{\pi(y)}\right) \pi(y) \geq \frac{\Psi(A)^2}{\pi(A^c)}.$$

Adding the upper and lower bounds for the sums over B and B^c gives bounds on $Q_{PP^*}(A, A^c)$:

$$\Psi(A)(2 - \tilde{\phi}(A)) \geq Q_{PP^*}(A, A^c) \geq \frac{\Psi(A)^2}{\pi(A)\pi(A^c)}$$

Dividing through by $\pi(A)\pi(A^c)$ and then re-arranging the inequalities completes the proof. \square

The lemma induces mixing bounds in terms of $\tilde{\Phi}_{PP^*}(r)$ for total variation, relative entropy and L^2 distance. For instance,

$$\tau_2(\epsilon) \leq \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{8 dr}{r \tilde{\Phi}_{PP^*}(r)^2} \right\rceil \quad \text{and} \quad \tau_2(\epsilon) \leq \left\lceil \frac{8}{\tilde{\Phi}_{PP^*}^2} \log \frac{1}{\epsilon \sqrt{\pi_*}} \right\rceil \quad (4.2)$$

This is not directly comparable to the Spectral profile bound, but it is never more than a factor two worse, and is strictly better when $x \tilde{\Phi}_{PP^*} \left(\frac{1}{1+x^2}\right)$ is convex as is often the case.

In a survey with Tetali [23] we use a more specialized method based on an idea of [25], which is applicable only to $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$, to show that

$$1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 1 - \sqrt[4]{1 - \Phi_{PP^*}(A)^2} \geq \frac{1}{4} \Phi_{PP^*}(A)^2.$$

This gives exactly the same mixing bound as the Spectral Profile result, and can be improved by a factor two when $x \Phi_{PP^*}^2 \left(\frac{1}{1+x^2}\right)$ is convex.

Now, consider the second bound of (4.1), with a holding probability. Modified conductance via Theorem 3.4.1, combined with Lemma 3.4.2, gives a weak bound on f -congestion in terms of conductance for non-lazy walks. We now give a more direct argument improving substantially on this.

Lemma 4.2.2. *Consider a Markov chain with holding probability $\gamma \in [0, 1]$. If $A \subset V$ then*

$$\begin{aligned} \tilde{\Phi}(A) &\geq 1 - \mathcal{C}_{a(1-a)}(A) &\geq \left(\frac{2}{\max\{\gamma, 1-\gamma\}} \right) \frac{\gamma}{1-\gamma} \tilde{\Phi}(A)^2 \pi(A) \pi(A^c) \\ \tilde{\Phi}(A) &\geq 1 - \mathcal{C}_{a \log(1/a)}(A) &\geq \left(\frac{1}{\max\{\gamma, 1-\gamma\}} \right) \frac{\gamma}{1-\gamma} \frac{\Phi(A)^2}{\log(1/\pi(A))} \\ \tilde{\Phi}(A) &\geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) &\geq \left(\frac{1}{4 \max\{\gamma, 1-\gamma\}} \right) \frac{\gamma}{1-\gamma} \tilde{\Phi}(A)^2 \end{aligned}$$

Proof. The upper bounds follow from Theorem 3.4.1 and the relation $\tilde{\phi}(A) \leq \tilde{\Phi}(A)$.

The lower bounds will be shown by using Lemma 3.6.2. The lazy walk P' has $\tilde{\phi}_{P'}(A) = \tilde{\Phi}_{P'}(A) = \frac{1}{2(1-\gamma)}\tilde{\Phi}(A)$.

The lower bounds with $\gamma > 1/2$ follow immediately from Lemma 3.6.2 and Theorem 3.4.1.

Now, the lower bound for $\gamma < 1/2$. By Theorem 3.1.2,

$$\frac{\mathbf{Q}(A, A^c)}{2(1-\gamma)} = \mathbf{Q}_{P'}(A, A^c) = \int_0^{1/2} (\pi(A'_w) - \pi(A)) dw = \int_{1/2}^1 (\pi(A) - \pi(A'_w)) dw.$$

Let $h(w) = \pi(A) + \frac{\mathbf{Q}(A, A^c)}{1-\gamma}$ if $w \leq 1/2$, and $h(w) = \pi(A) - \frac{\mathbf{Q}(A, A^c)}{1-\gamma}$ if $w > 1/2$. This satisfies the conditions of Lemma 3.6.2. Theorem 3.4.1 then completes the proof, for instance if $f(a) = \sqrt{a(1-a)}$ then

$$1 - \mathcal{C}_f(A) \geq 2\gamma \left(1 - \frac{\int_0^1 f(h(w)) dw}{f(\pi(A))} \right) \geq 2\gamma \frac{\tilde{\Phi}_{P'}(A)^2}{2} = \frac{\gamma}{4(1-\gamma)^2} \tilde{\Phi}(A)^2.$$

□

The lemma induces mixing bounds in terms of $\tilde{\Phi}(r)$ for total variation, relative entropy and L^2 distance. For instance,

$$\tau_2(\epsilon) \leq \left\lceil \int_{4\pi_*}^{4/\epsilon^2} \frac{4 \max\{\gamma, 1-\gamma\}}{\frac{\gamma}{1-\gamma} r \tilde{\Phi}(r)^2} dr \right\rceil \quad \text{and} \quad \tau_2(\epsilon) \leq \left\lceil \frac{4 \max\{\gamma, 1-\gamma\}}{\frac{\gamma}{1-\gamma} \tilde{\Phi}^2} \log \frac{1}{\epsilon \sqrt{\pi_*}} \right\rceil. \quad (4.3)$$

This is not directly comparable to the Spectral profile bound, but it is never more than a factor two worse, and is strictly better when the walk is lazy (i.e. $\gamma = 1/2$) or $x\tilde{\Phi}^2\left(\frac{1}{1+x^2}\right)$ is convex.

4.3 Blocking Conductance

As discussed in the introduction, our methods give new insight into the mixing time bounds of Blocking conductance [13]. We note that the improved Average Conductance result of Fountoulakis and Reed [8] is actually a special case of the Blocking Conductance total variation bound discussed below, so our discussion applies to their results as well.

In this section we work with the quantity $\Psi(A, t)$, first defined in equation (3.2). That definition was only used for $t \leq \pi(A^c)$, but when $t > \pi(A^c)$ a different definition is more appropriate. In general, let

$$\Psi(A, t) = \min_{\substack{B \subset V, v \in V, \\ \pi(B) \leq t, \pi(B \cup v) > t}} \mathbf{Q}(A, B) + (t - \pi(B)) \frac{\mathbf{Q}(A, v)}{\pi(v)} \quad (4.4)$$

if $t \leq \pi(A^c)$ and $\Psi(A, t) = \Psi(A^c, 1-t)$ when $t > \pi(A^c)$.

The Blocking Conductance theorem is the following:

Theorem 4.3.1. [Blocking Conductance [13]] *Given a lazy, reversible, ergodic Markov chain then*

$$\tau_{TV}(\epsilon) \leq 15000 \left(\int_{\pi_*}^{1/2} \hat{h}(x) dx + \hat{h}(1/2) \right) \log_2(1/2\epsilon)$$

where $\hat{h}(x)$ can be any of the following:

1. $\forall x \in [0, 1] : \hat{h}_{gl}(x) \geq \sup_{\substack{A \subset \hat{V}, \\ \pi(A) \leq x}} \frac{1}{\pi(A) \psi_{gl}(A)}$ where $\psi_{gl}(A) = \int_0^1 \frac{\Psi(A^c, t)}{\pi(A)^2} dt$
2. $\forall x \in [0, 1] : \hat{h}_{mod}(x) \geq \sup_{\substack{A \subset \hat{V}, \\ \pi(A) \leq x}} \frac{1}{x \psi_{mod}(A)}$ where $\psi_{mod}(A) = \int_0^1 \frac{\Psi(A^c, t)}{t \pi(A)} dt$
3. $\forall x \in [0, 1] : \hat{h}^+(x) \geq \sup_{\substack{A \subset \hat{V}, \\ x/2 \leq \pi(A) \leq x}} \frac{1}{x \psi^+(A)}$ where $\psi^+(A) = \int_0^{\pi(A)} \frac{\Psi(A^c, t)}{\pi(A)^2} dt$

The state space $\hat{V} = [0, 1]$ is the continuization of V , and is defined by associating to each $v \in V$ a disjoint interval of size $\pi(v)$, with ergodic flow such that if $dx \subset v_1$ and $dy \subset v_2$ then $\mathbf{Q}(dx, dy) = dx \mathbf{P}(v_1, v_2) \frac{dy}{\pi(v_2)}$.

The large coefficient is due to a conversion from one measure of mixing time to another, and the need for the continuization is because the theorem is proven in the continuous space setting. A discrete version of this is discussed in Section 4.5.

To relate this to Evolving sets we first rewrite $\Psi(A, t)$ in terms of Evolving Sets [20]:

Lemma 4.3.2. *Given a finite irreducible Markov kernel and $A \subset V$ then*

$$\Psi(A^c, t) = \begin{cases} \int_{w(t)}^1 (t - \pi(A_u)) du & \text{if } t \leq \pi(A) \\ \int_0^{w(t)} (\pi(A_u) - t) du & \text{if } t \geq \pi(A) \end{cases}$$

where $w(t)$ is any value satisfying $\inf\{y : \pi(A_y) \leq t\} \leq w(t) \leq \sup\{y : \pi(A_y) \geq t\}$.

Proof. We consider only the case that $t \leq \pi(A)$ since the case when $t > \pi(A)$ is similar.

By definition, if $v_1 \in A_u$ and $v_2 \notin A_u$ then $\frac{\mathbf{Q}(A, v_1)}{\pi(v_1)} > \frac{\mathbf{Q}(A, v_2)}{\pi(v_2)}$, and equivalently $\frac{\mathbf{Q}(A^c, v_1)}{\pi(v_1)} < \frac{\mathbf{Q}(A^c, v_2)}{\pi(v_2)}$. Hence, if $\pi(A_{w(t)}) = t$ then $B = A_{w(t)}$ is the same set where the minimum occurs in the definition of $\Psi(A^c, t)$. If instead $\pi(A_{w(t)}) > t$, then $B = \cup_{u > w(t)} A_u$ is the set where the minimum occurs in the definition of $\Psi(A^c, t)$, and if v is any vertex in $A_{w(t)} \setminus B$ then $\mathbf{Q}(A, v)/\pi(v) = w(t)$. In both cases

$$\Psi(A^c, t) = \mathbf{Q}(A^c, B) + (t - \pi(B)) \frac{\mathbf{Q}(A^c, v)}{\pi(v)}.$$

Let B be as defined in the previous paragraph. Then, $B \subseteq A_{w(t)}$, and $A_u \subseteq B$ whenever $u > w(t)$, and so

$$\begin{aligned} \int_{w(t)}^1 (t - \pi(A_u)) du &= t(1 - w(t)) - \sum_{y \in B} \left(\frac{\mathbf{Q}(A, y)}{\pi(y)} - w(t) \right) \pi(y) \\ &= t(1 - w(t)) - (\mathbf{Q}(A, B) - w(t)\pi(B)) \\ &= (1 - w(t))(t - \pi(B)) + \mathbf{Q}(A^c, B) \\ &= \mathbf{Q}(A^c, B) + (t - \pi(B)) \frac{\mathbf{Q}(A^c, v)}{\pi(v)} = \Psi(A^c, t). \end{aligned}$$

The first equality is because $\int_x^1 \pi(A_u) du = \sum_{y \in A_x} (\text{Prob}(y \in A_u) - x)\pi(y)$. The third equality uses $Q(A, B) = \pi(B) - Q(A^c, B)$. The fourth equality is because $\frac{Q(A^c, v)}{\pi(v)} = 1 - \frac{Q(A, v)}{\pi(v)} = 1 - w(t)$ by our choice of v and $w(t)$. \square

The next step is to rewrite f -congestion quantities in terms of the $\psi(A)$ quantities appearing in the Blocking Conductance theorem.

Lemma 4.3.3. *Let $\tilde{\psi}_{gl}(A) = \int_0^1 \frac{\Psi(A^c, t)}{\pi(A)^2 \pi(A^c)^2} dt$ and $\tilde{\psi}^+(A) = \int_0^{\pi(A)} \frac{\Psi(A^c, t)}{\pi(A)^2 \pi(A^c)^2} dt$. Then,*

$$\begin{aligned} 1 - \mathcal{C}_{a(1-a)}(A) &= 2\pi(A)\pi(A^c)\tilde{\psi}_{gl}(A) \\ 1 - \mathcal{C}_{a \log(1/a)}(A) &= \frac{\psi_{mod}(A)}{\log \frac{1}{\pi(A)}} \\ 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) &\geq \frac{1}{4} \tilde{\psi}^+(A). \end{aligned}$$

The identity $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A^c)$ can be used when $\pi(A) > 1/2$.

Proof. Start with the case of $1 - \mathcal{C}_{a \log(1/a)}(A)$.

First, rewrite things a bit.

$$\begin{aligned} 1 - \mathcal{C}_{a \log(1/a)}(A) &= \frac{\pi(A) \log(1/\pi(A)) - \int_0^1 \pi(A_u) \log(1/\pi(A_u)) du}{\pi(A) \log(1/\pi(A))} \\ &= \frac{1}{\log \frac{1}{\pi(A)}} \int_0^1 \int_{\pi(A_u)}^{\pi(A)} \frac{t - \pi(A_u)}{t \pi(A)} dt du \end{aligned}$$

The second equality applied the identity $\int_0^1 \pi(A_u) du = \pi(A)$.

Now to rewrite $\psi_{mod}(A)$ in terms of Evolving Sets. By Lemma 4.3.2,

$$\begin{aligned} \int_0^1 \frac{\Psi(A^c, t)}{t \pi(A)} dt &= \int_0^{\pi(A)} \int_{w(t)}^1 \frac{t - \pi(A_u)}{t \pi(A)} du dt + \int_{\pi(A)}^1 \int_0^{w(t)} \frac{\pi(A_u) - t}{t \pi(A)} du dt \\ &= \int_0^1 \int_{\pi(A_u)}^{\pi(A)} \frac{t - \pi(A_u)}{t \pi(A)} dt du \end{aligned} \tag{4.5}$$

The $1 - \mathcal{C}_{a(1-a)}(A)$ result is shown similarly.

For the $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ case we first re-write $\tilde{\psi}(A)^+$ in terms of Evolving Sets. Let $\wp_A = w(\pi(A))$, as in previous sections. Then

$$\begin{aligned} \int_0^{\pi(A)} \frac{\Psi(A^c, t)}{\pi(A)^2 \pi(A^c)^2} dt &= \int_0^{\pi(A)} \int_{w(t)}^1 \frac{t - \pi(A_u)}{\pi(A)^2 \pi(A^c)^2} du dt \\ &= \int_{\wp_A}^1 \int_{\pi(A_u)}^{\pi(A)} \frac{t - \pi(A_u)}{\pi(A)^2 \pi(A^c)^2} dt du \\ &= \frac{1}{2} \int_{\wp_A}^1 \frac{(\pi(A) - \pi(A_u))^2}{\pi(A)^2 \pi(A^c)^2} du \end{aligned} \tag{4.6}$$

Finish with the inequality that for $x, y \in [0, 1]$

$$\sqrt{y(1-y)} \leq \sqrt{x(1-x)} + \frac{1-2x}{2\sqrt{x(1-x)}}(y-x) - \frac{(y-x)^2}{2[x(1-x)]^{3/2}} \left(\delta_{y \leq x} \frac{1}{4} + \delta_{y > x} x(1-x) \right),$$

substitute $y = \pi(A_u)$ and $x = \pi(A)$, and then integrate to obtain $\mathcal{C}_{\sqrt{a(1-a)}}(A) \leq 1 - \frac{1}{4} \tilde{\psi}^+(A)$. The relation follows from the inequality $\sqrt{z} \leq 1 + \frac{1}{2}(z-1) - \frac{1}{8}(z-1)^2 \delta_{z \leq 1}$ with $z = \frac{y(1-y)}{x(1-x)}$. \square

When combined with Lemma 2.3.2 it follows, for instance, that

$$\hat{\mathbb{E}}_{n+1} \log \frac{1}{\pi(S_{n+1})} - \hat{\mathbb{E}}_n \log \frac{1}{\pi(S_n)} = -\hat{\mathbb{E}}_n \psi_{mod}(S_n).$$

Hence the expectation of $\psi_{mod}(A)$ is exactly the rate at which the evolving set bound on relative entropy decreases. This shows that in a sense Blocking Conductance and Evolving Set bounds are both based on measuring the derivative of the distance with respect to time. Not surprisingly, the Evolving set mixing bounds then imply bounds of the Blocking Conductance form.

Corollary 4.3.4. *Consider a finite (non-lazy, non-reversible) ergodic Markov chain. Then*

$$\begin{aligned} \tau_{TV}(\epsilon) &\leq \left\lceil \frac{1}{2} h_{gl}(1/2) \log \frac{1 - \pi_*}{\epsilon} \right\rceil \\ \tau_D(\epsilon) &\leq \left\lceil 2C \int_{\sqrt{\pi_*}}^{1/2} h_{mod}(x) dx + C h_{mod}(1/2) \log \frac{2}{\epsilon} \right\rceil \\ \tau_2(\epsilon) &\leq \left\lceil 4 \int_{4\pi_*}^{1/2} h^+(x) dx + h^+(1/2) \log \frac{2\sqrt{2}}{\epsilon} \right\rceil \end{aligned}$$

where

$$h_{gl}(x) = \max_{\substack{A \subset V, \\ \pi(A) \leq x}} \frac{1}{\pi(A) \psi_{gl}(A)}, \quad h_{mod}(x) = \max_{\substack{A \subset V, \\ \pi(A) \leq x}} \frac{1}{x \psi_{mod}(A)}, \quad h^+(x) = \max_{\substack{A \subset V, \\ \pi(A) \leq x}} \frac{1}{x \psi^+(A)}$$

and C is the optimal constant satisfying

$$\forall r \geq \pi_* : \min_{\pi(A) \leq r} \frac{\psi_{mod}(A)}{\log(1/\pi(A))} \geq C^{-1} \min_{\pi(A) \leq r} \frac{\psi_{mod}(A)}{\log(1/r)}.$$

Proof. For the total variation and L^2 bounds apply Corollary 2.3.3 and Theorem 2.3.4 respectively to obtain mixing time bounds in terms of various $1 - \mathcal{C}_f(A)$. Replacing the f -congestion by the appropriate $\psi(A)$ quantities from Lemma 4.3.3 then gives the results. However, the relative entropy case requires more work. This is because $\forall r > 1/2$ both $1 - \mathcal{C}_{a(1-a)}(r) = 1 - \mathcal{C}_{a(1-a)}(1/2)$ and $1 - \mathcal{C}_{\sqrt{a(1-a)}}(r) = 1 - \mathcal{C}_{\sqrt{a(1-a)}}(1/2)$, while $1 - \mathcal{C}_{a \log(1/a)}(r) \neq 1 - \mathcal{C}_{a \log(1/a)}(1/2)$ when $r > 1/2$.

From Theorem 2.2.2, it follows that if $g(a) = \min\{1 + \log \frac{1}{2a}, \frac{1-a}{a}(1 + \log \frac{1}{2(1-a)})\}$ then

$$D(\mathbb{P}^n(x, \cdot) \| \pi) \leq \hat{\mathbb{E}}_n \log \frac{1}{\pi(S_n)} \leq \hat{\mathbb{E}}_n g(\pi(S_n)).$$

By Theorem 2.3.4, and the relation $\mathcal{C}_{ag(a)}(r) = \mathcal{C}_{ag(a)}(1/2)$ for $r \geq 1/2$ (since $ag(a) = (1-a)g(1-a)$), the mixing time is then bounded by

$$\tau_D(\epsilon) \leq \left[\int_{\sqrt{e\pi^*/2}}^{1/2} \frac{2 dr}{r(1 + \log(1/2r))(1 - \mathcal{C}_{ag(a)}(r))} + \frac{2 \log \frac{2}{\epsilon}}{1 - \mathcal{C}_{ag(a)}(1/2)} \right].$$

Consider set $A \subset V$ with $\pi(A) \leq r \leq 1/2$. Then $\pi(A)g(\pi(A)) = \pi(A)(1 + \log \frac{1}{2\pi(A)})$ and $ag(a) \leq a(1 + \log \frac{1}{2a}) \forall a \in [0, 1]$, and so

$$\begin{aligned} 1 - \mathcal{C}_{ag(a)}(A) &\geq 1 - \mathcal{C}_{a(1+\log(1/2a))}(A) \\ &= \frac{\log(1/\pi(A))}{1 + \log(1/2\pi(A))} (1 - \mathcal{C}_{a \log(1/a)}(A)) \\ &\geq \frac{\log(1/r)}{1 + \log(1/2r)} C^{-1} \frac{\psi_{mod}(r)}{\log(1/r)}. \end{aligned}$$

Substituting this into the bound on $\tau_D(\epsilon)$ given above completes the proof. \square

The Corollary shows that as long as the bottlenecks get sufficiently worse as set size increases, then Evolving set bounds sharply improve on Blocking conductance results. To see this for the h_{mod} case note that if $\min_{\pi(A) \leq r} \frac{\psi_{mod}(A)}{\log(1/\pi(A))} = \min_{\pi(A) \in [r/2, r]} \frac{\psi_{mod}(A)}{\log(1/\pi(A))}$, i.e. congestion decreases with set size, then it suffices to take $C = 2$. The new bounds are, however, usually much better than the Blocking Conductance bounds because the laziness and reversibility requirements have been dropped, the bounds are given in terms of stronger measures of distance, and there is no need to work in a continuous state space.

This suggests that the Blocking Conductance method is the best method only in the case when bottlenecks are worst at small sets. This situation can arise when total variation mixing time is an order of magnitude faster than relative-entropy or L^2 -mixing times. For instance, Fountoulakis and Reed [8] use a version of Blocking Conductance to find the correct mixing time of walks on the giant component of the random graph $G_{n,p}$. The bottleneck condition also does not apply to certain walks used for estimating volume of convex bodies, or to Example 4.3.7 below.

The interested reader can use the quantities calculated in Example 3.7.3 to find that Corollary 4.3.4 is within a factor 4 of being sharp for the walk on a complete graph. A ‘‘convex’’ version, based on Theorem 2.3.4, can be used to strengthen this to a factor 2.

Remark 4.3.5. *A straightforward generalization of work in [20] can be used to show that*

$$\psi_{gl}(A) \geq \frac{1}{2} \psi_{mod}(A) \geq 1 - \mathcal{C}_{\sqrt{a}}(A) \geq \frac{1}{4} \psi^+(A) \geq \frac{1}{4} \phi(A)^2.$$

Hence, these various $\psi(A)$ quantities are closely related to each other, and to modified conductance.

Remark 4.3.6. *For a lazy walk a useful interpretation of $\psi^+(A)$ is given in [13]:*

$$\psi^+(A) \geq \sup_{\lambda \leq \pi(A)} \min_{\substack{S \subset A, \\ \pi(S) < \lambda}} \frac{\lambda \mathbb{Q}(A \setminus S, A^c)}{\pi(A)^2} \geq \frac{1}{2} \Phi^2(A).$$

When combined with Lemma 4.3.3 it follows that

$$1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq \frac{1}{4} \sup_{\lambda \leq \pi(A)} \min_{\substack{S \subset A, \\ \pi(S) < \lambda}} \frac{\lambda \mathbf{Q}(A \setminus S, A^c)}{\pi(A)^2 \pi(A^c)^2} \geq \frac{1}{8} \tilde{\Phi}^2(A).$$

This can be interpreted as follows. Let λ denote the maximal size of a “blocking set”, such that if any set S smaller than this is blocked from transitioning then it does not block too much of the ergodic flow $\mathbf{Q}(A, A^c)$. For instance, $\mathbf{Q}(A \setminus S, A^c) = \mathbf{Q}(A, A^c) - \mathbf{Q}(S, A^c) \geq \mathbf{Q}(A, A^c) - \lambda/2$, and so by setting $\lambda = \mathbf{Q}(A, A^c)$ then the first lower bound on $\psi^+(A)$ implies the second.

See Remark 4.6.3 for a similar discussion involving the non-lazy case.

Example 4.3.7. The L^2 mixing time can be slower than total variation mixing time when there is a bottleneck at a small set, in which case the difference between Theorem 4.3.1 and Corollary 4.3.4 may be real, and not simply an artifact of the method of proof.

Consider the complete graph K_m on m vertices and attach an additional vertex v by a single edge. We examine the lazy max-degree walk given by choosing a neighboring vertex with probability $1/2m$ each, and otherwise do nothing.

First, bound $\psi^+(A)$. If $A = \{v\}$ then let $\lambda = \pi(\{v\}) = \frac{1}{m+1}$. The only set $\pi(S) < \lambda$ is $S = \emptyset$, and so $\psi^+(\{v\}) \geq \frac{\lambda \mathbf{Q}(\{v\}, K_m)}{\pi(\{v\})^2} = \frac{1}{2m}$. If $A \neq \{v\}$ then $\Phi(A) \geq 1/8$, and so $\psi^+(A) \geq \frac{1}{128}$.

To bound mixing via Blocking Conductance, note that $\hat{h}^+(r) \leq \frac{2m}{x}$ if $r \leq \frac{1}{2(m+1)}$, while $\hat{h}^+(r) \leq \frac{128}{x}$ otherwise. Then, by Theorem 4.3.1,

$$\tau_{TV}(\epsilon) = O(m \log(1/\epsilon)),$$

which is of the correct order.

For Evolving Sets, we can only say that $h^+(r) \leq \frac{2m}{x}$ for all r . Then, by Corollary 4.3.4,

$$\tau_2(\epsilon) = O(m \log(m/\epsilon))$$

which is again of the correct order.

4.4 Comparison and Canonical paths for Evolving Sets

One of the most useful isoperimetric arguments for bounding mixing times has been the method of Canonical Paths, and in its more general form the method of Comparison. In this section we find similar results for Evolving Sets, although our results are somewhat weaker than might be hoped as they only allow us to compare the f -congestion \mathcal{C}_f of one Markov chain with the Conductance Profile $\tilde{\Phi}(r)$ of another. Nevertheless, the results do serve to provide an overarching framework explaining why different versions of canonical path bounds are stronger in different situations, and in particular why ρ_e (edge-congestion) can be multiplied by any of ρ_e (edge-congestion), or ρ_v (vertex-congestion) or ℓ (maximum path length) in order to bound spectral gap (and mixing time): $\lambda \geq c/\rho_e \rho_e, c/\rho_e \rho_v, c/\rho_e \ell$. Certainly ρ_e is a measure of edge-congestion, but it is perhaps more surprising that the terms potentially multiplied by this are all measures of vertex-congestion; for a lazy chain $\rho_v \geq \rho_e/2$ and so ρ_e bounds vertex congestion, certainly ρ_v does as well, and we will find that ℓ is a guarantor of good “typical” vertex congestion in some well defined sense.

In [21] we found that if only the worst case edge expansion (i.e., conductance or edge congestion) are known then the best evolving set bound that can be obtained for spectral gap is via $1 - \mathcal{C}_{\sqrt{a(1-a)}}$, which we have seen to be a natural quantity for bounding L^2 or chi-square distance. The same idea suggests this to be the best that can be done for mixing time as well. This partially explains why all path bounds have been bounds on chi-square distance, since paths are generally used to show good edge expansion.

We start with an elementary comparison result, comparison of ergodic flows, to illustrate the use of canonical paths.

Theorem 4.4.1 (Comparison of Flows). *Suppose \mathcal{M} and \mathcal{M}' are finite Markov chains on state space V and edge sets E and E' respectively. To every edge $e' = (x, y) \in E'$ associate a path $\gamma_{xy} \subset E$ from x to y . Let*

$$\rho_e = \max_{e=(u,v) \in E} \frac{1}{\pi(u)P(u,v)} \sum_{\gamma_{xy} \ni e} \pi'(x) P'(x,y).$$

Then, for every $A \subset V$,

$$Q(A, A^c) \geq Q'(A, A^c) / \rho_e.$$

Proof. For each edge $e' = (x, y) \in E'$ with $x \in A$ and $y \in A^c$, transport flow of $\pi'(x) P'(x, y)$ along the path $\gamma_{xy} \subset E$ from x to y , for a total of $Q'(A, A^c)$ from A to A^c . By definition of ρ_e it follows that if $e = (u, v) \in E$ with $u \in A$, $v \in A^c$ then $\pi(u)P(u, v) \geq \rho_e^{-1} \sum_{\gamma_{xy} \ni e} \pi'(x) P'(x, y)$, and therefore

$$Q(A, A^c) = \sum_{\substack{e=(u,v) \in E, \\ u \in A, v \in A^c}} \pi(u)P(u, v) \geq \frac{\sum_{x \in A, y \in A^c} \pi'(x) P'(x, y)}{\rho_e} = \frac{Q'(A, A^c)}{\rho_e}$$

because if $(x, y) \in A \times A^c$ then γ_{xy} must include some edge $(u, v) \in A \times A^c$. \square

It follows that $\tilde{\Phi}(A) \geq \rho_e^{-1} \frac{Q'(A, A^c)}{\pi(A)\pi(A^c)}$, and in particular when $\pi = \pi'$ then $\tilde{\Phi}(A) \geq \tilde{\Phi}'(A) / \rho_e$. By equation (3.1), if \mathcal{M} is lazy and $\pi = \pi'$ then $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 1 - \sqrt{1 - (\tilde{\Phi}'(A) / \rho_e)^2}$. Furthermore, since spectral gap determines the asymptotic rate of convergence, then also if \mathcal{M} is reversible then

$$\lambda \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}} \geq 1 - \sqrt{1 - (\tilde{\Phi}' / \rho_e)^2} \geq \tilde{\Phi}'^2 / 2\rho_e^2. \quad (4.7)$$

As with comparison methods, by comparing to the complete graph we obtain a “standard” canonical path type bound. In this case let \mathcal{M}' have transitions $P'(x, y) = \pi(y)$, so that $Q'(A, A^c) = \pi(A)\pi(A^c)$ and $\tilde{\Phi}'(A) = 1$. Then

$$\lambda \geq 1 - \sqrt{1 - 1/\rho_e^2} \geq 1/2\rho_e^2,$$

a factor of 4 improvement over Jerrum and Sinclair’s bound $\lambda \geq 1/8\rho_e^2$ [11, 27].

Our main result in this section is a comparison theorem in terms of edge and vertex congestion. This applies to a wide range of distance measures and to non-reversible Markov chains.

Theorem 4.4.2. *Suppose \mathcal{M} and \mathcal{M}' are finite Markov chains on state space V and edge sets E and E' respectively. To every edge $e' = (x, y) \in E'$ associate a path $\gamma_{xy} \subset E$. Let*

$$\rho_e = \max_{e \in E} \frac{1}{\mathbf{Q}(e)} \sum_{\gamma_{xy} \ni e} \pi'(x) \mathbf{P}'(x, y) \quad \text{and} \quad \rho_v = \max_{v \in V} \frac{1}{\pi(v)} \sum_{\gamma_{xy} \ni v} \pi'(x) \mathbf{P}'(x, y).$$

If \mathcal{M} is lazy, $A \subset V$, and $f : [0, 1] \rightarrow \mathbb{R}$ is a concave function with $f(\pi(A)) \neq 0$ then

$$1 - \mathcal{C}_f(A) \geq 2 \frac{\rho_v}{\rho_e} \left(1 - \frac{f\left(\pi(A) + \frac{\mathbf{Q}'(A, A^c)}{\rho_v}\right) + f\left(\pi(A) - \frac{\mathbf{Q}'(A, A^c)}{\rho_v}\right)}{2f(\pi(A))} \right).$$

In the sum for ρ_v if both paths γ_{xy} and γ_{yx} pass through the vertex v then it suffices to replace the sum of weights $\pi'(x) \mathbf{P}'(x, y) + \pi'(y) \mathbf{P}'(y, x)$ by $\max\{\pi'(x) \mathbf{P}'(x, y), \pi'(y) \mathbf{P}'(y, x)\}$.

Proof. Given a set $A \subset V$, by Lemma 3.3.1 to lower bound $1 - \mathcal{C}_f(A)$ it suffices to construct the distribution of $\pi(A_u)$ which minimizes $\int_0^t \pi(A_u) du$.

Recall that $\int_0^{1/2} \pi(A_u) du = \frac{\pi(A)}{2} + \mathbf{Q}(A, A^c)$. Therefore, given the correct distribution of $\pi(A_u)$ on $[0, 1/2]$, and an underestimate of the ergodic flow $\mathbf{Q}_0(A, A^c) \leq \mathbf{Q}(A, A^c)$, then the distribution $\pi_0(A_u) = \pi(A) + (\pi(A_u) - \pi(A)) \frac{\mathbf{Q}_0(A, A^c)}{\mathbf{Q}(A, A^c)}$ will certainly satisfy $\forall t \in [0, 1/2] : \int_0^t \pi_0(A_u) du \leq \int_0^t \pi(A_u) du$ and $\int_0^{1/2} \pi_0(A_u) du = \frac{\pi(A)}{2} + \mathbf{Q}_0(A, A^c)$. A similar argument holds for $t \in [1/2, 1]$. It follows that underestimating the ergodic flow by only considering that due to the paths will only make the lower bound on $1 - \mathcal{C}_f(A)$ too pessimistic.

Now, let \mathbf{Q}_0 denote the ergodic flow measured by the paths, so that $\forall v \in A^c : \mathbf{Q}_0(A, v) = \sum_{x \in A, y \in A^c} \delta_{\gamma_{xy} \ni v} \pi'(x) \mathbf{P}'(x, y) / \rho_e$ and likewise for $v \in A$, so $\mathbf{Q}_0(A, A^c) \geq \mathbf{Q}'(A, A^c) / \rho_e$. The vertex congestion implies that not too much of this passes through any specific vertex, and in particular $\forall v \in A : \frac{\mathbf{Q}_0(A^c, v)}{\pi(v)} \leq \frac{\mathbf{Q}'(A, A^c) / \rho_e}{\mathbf{Q}'(A, A^c) / \rho_v} = \frac{\rho_v}{\rho_e}$, while $\forall v \in A^c : \frac{\mathbf{Q}_0(A, v)}{\pi(v)} \leq \frac{\rho_v}{\rho_e}$ as well.

It follows that if $M := \rho_v / \rho_e$ then $\forall u \in [M, 1 - M] : \pi(A_u) = \pi(A)$. Subject to this constraint, in Lemma 3.3.1 the integral $\int_0^t \pi(A_u) du$ is minimized by

$$\pi(A_u) = \begin{cases} \pi(A) + M^{-1} \mathbf{Q}(A, A^c) & \text{if } u < M \\ \pi(A) & \text{if } u \in [M, 1 - M] \\ \pi(A) - M^{-1} \mathbf{Q}(A, A^c) & \text{if } u > 1 - M \end{cases}$$

For these values of $\pi(A_u)$ integration shows that

$$\int_0^1 f(\pi(A_u)) du = M f\left(\pi(A) + \frac{\mathbf{Q}(A, A^c)}{M}\right) + (1 - 2M) f(\pi(A)) + M f\left(\pi(A) - \frac{\mathbf{Q}(A, A^c)}{M}\right)$$

By definition $M = \rho_v / \rho_e$, and by the remarks above $\mathbf{Q}(A, A^c) = \mathbf{Q}'(A, A^c) / \rho_e$, which combined with the integral above gives the result.

The final comment on ρ_v follows because if $v \in A$ then we need only consider congestion among paths entering A , and as only one of γ_{xy} or γ_{yx} will enter A then there is no need to count congestion due to both paths. Likewise if $v \in A^c$. \square

This can be used to show lower bounds for the various f -congestion quantities, and hence upper bounds on each notion of mixing times.

Corollary 4.4.3 (Comparison with Conductance Function). *Suppose \mathcal{M} and \mathcal{M}' are finite Markov chains on state space V and edge sets E and E' respectively. If $A \subset V$ is a proper subset and $\pi = \pi'$ then*

$$\begin{aligned} 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) &\geq 2 \frac{\rho_v}{\rho_e} \left(1 - \sqrt{1 - \left(\frac{\tilde{\Phi}'(A)}{2\rho_v} \right)^2} \right) \geq \frac{1}{4} \frac{\tilde{\Phi}'(A)^2}{\rho_v \rho_e} \\ 1 - \mathcal{C}_{a \log(1/a)}(A) &\geq \frac{1}{\log(1/\pi(A))} \frac{\Phi'(A)^2}{\rho_v \rho_e} \\ 1 - \mathcal{C}_{a(1-a)}(A) &\geq 2\pi(A)\pi(A^c) \frac{\tilde{\Phi}'(A)^2}{\rho_v \rho_e} \end{aligned}$$

If $\pi \neq \pi'$ then replace $\tilde{\Phi}'(A)$ by $\frac{\mathcal{Q}'(A, A^c)}{\pi(A)\pi(A^c)}$ and likewise $\Phi'(A)$ by $\frac{\mathcal{Q}'(A, A^c)}{\pi(A)}$.

Proof. Using the previous lemma it is just a matter of simplification.

We will save ourselves the work of simplifying by instead reducing this to a problem that was already solved previously. Observe that if

$$\pi(A_u) = \begin{cases} \pi(A) - \frac{\mathcal{Q}'(A, A^c)/2\rho_v}{1/2} & \text{if } u > 1/2 \\ \pi(A) + \frac{\mathcal{Q}'(A, A^c)/2\rho_v}{1/2} & \text{if } u \leq 1/2 \end{cases} \quad (4.8)$$

then

$$1 - \mathcal{C}_f(A) = 1 - \frac{f\left(\pi(A) + \frac{\mathcal{Q}'(A, A^c)}{\rho_v}\right) + f\left(\pi(A) - \frac{\mathcal{Q}'(A, A^c)}{\rho_v}\right)}{2f(\pi(A))}$$

is exactly the same as the lower bound on $\frac{1 - \mathcal{C}_f(A)}{2\rho_v/\rho_e}$ given in Theorem 4.4.2.

However, the distribution (4.8) is the worst case bound for $\pi(A_u)$ for a lazy chain with ergodic flows $\mathcal{Q}(A, A^c) \geq \mathcal{Q}'(A, A^c)/2\rho_v$ (see Figure 3.1). We already determined that this leads to the bound $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 1 - \sqrt{1 - \tilde{\Phi}(A)^2}$. Therefore,

$$\frac{1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)}{2\rho_v/\rho_e} \geq 1 - \sqrt{1 - \left(\frac{\mathcal{Q}'(A, A^c)/2\rho_v}{\pi(A)\pi(A^c)} \right)^2}$$

and the lower bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ follows immediately. The other bounds follow similarly. \square

If $\mathcal{M}' = \mathcal{M}$, \mathcal{M} is reversible and $\forall x \in V : \mathbf{P}(x, x) \geq \gamma \in [1/2, 1]$ then $\rho_v \leq (1 - \gamma)\rho_e = 1 - \gamma$ and so

$$\tilde{\Phi}(A) \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq 2(1 - \gamma) \left(1 - \sqrt{1 - \left(\frac{\tilde{\Phi}(A)}{2(1 - \gamma)} \right)^2} \right) \geq \frac{\tilde{\Phi}(A)^2}{4(1 - \gamma)}.$$

Without reversibility we have only that $\rho_v \leq 2(1 - \gamma)\rho_e$, and a slightly weaker result is obtained.

Corollary 4.4.3 is all that is needed to prove Corollary 4.4.4. In the particular case when \mathcal{M}' has transition probabilities $\mathbf{P}'(x, y) = \pi(y)$ then $\tilde{\Phi}'(A) = 1$ and this reduces to a ‘‘canonical paths’’ theorem.

Corollary 4.4.4 (Canonical Paths). *Suppose \mathcal{M} is a finite ergodic lazy Markov chain on state space V and edge set E , and Γ is a set of paths γ_{xy} between every pair of distinct vertices $x, y \in V$. Define*

$$\rho_e = \max_{e \in E} \frac{1}{\mathbf{Q}(e)} \sum_{\gamma_{xy} \ni e} \pi(x) \pi(y) \quad \text{and} \quad \rho_v = \max_{v \in V} \frac{1}{\pi(v)} \sum_{\gamma_{xy} \ni v} \pi(x) \pi(y).$$

Then

$$1 - \mathcal{C}_{\sqrt{x(1-x)}} \geq 2 \frac{\rho_v}{\rho_e} \left(1 - \sqrt{1 - 1/4\rho_v^2}\right) \geq \frac{1}{4\rho_v \rho_e}$$

and

$$\tau(\epsilon) \leq 4\rho_v \rho_e \log \frac{1}{2\epsilon\sqrt{\pi_*}}.$$

If \mathcal{M} is reversible then ρ_e and ρ_v can be taken as sums over unordered pairs (x, y) with undirected edges, and moreover

$$\lambda \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}} \geq 2 \frac{\rho_v}{\rho_e} \left(1 - \sqrt{1 - 1/4\rho_v^2}\right) \geq \frac{1}{4\rho_v \rho_e},$$

where $\lambda = 1 - \lambda_2$ is the spectral gap between 1 and the second largest eigenvalue λ_2 of the transition matrix \mathbf{P} .

Proof. Let \mathcal{M}' be a walk on the state space V with transition probabilities $\mathbf{P}'(x, y) = \pi(y)$. Then $\pi' = \pi$, $\mathbf{Q}'(A, A^c) = \pi(A)\pi(A^c)$, to every edge in \mathcal{M}' (i.e. pair of vertices) is associated a path γ_{xy} given in the corollary, and the congestions are exactly the ρ_e and ρ_v given in the corollary. The bounds on λ , $1 - \mathcal{C}_{\sqrt{a(1-a)}}$ and $\tau(\epsilon)$ then follow from Corollary 4.4.3. \square

As mentioned earlier, for a lazy chain $\rho_v \leq \rho_e/2$, which reduces the result to $\lambda \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}} \geq 1 - \sqrt{1 - \rho_e^{-2}}$, which we already know is sharp for the lazy walk on the uniform two-point space. Our extension shows that the second edge congestion term in the canonical paths bound $\tau = O(\rho_e \cdot \rho_e \log \pi_0^{-1})$ [11, 27] should actually have been the smaller vertex congestion ρ_v . This can be a significant improvement, as the following example shows.

Example 4.4.5. Feder and Mihail [5] studied a random walk for sampling balanced matroids and showed a result equivalent to $\rho_e \leq nm$ and $\rho_v \leq 2n$. The Cheeger bound (Jerrum and Sinclair [11]) implies a bound of $\tau(\epsilon) \leq 8\rho_e^2 \log \frac{1}{2\epsilon\sqrt{\pi_*}} \leq 8m^2 n^2 \left(\frac{n}{2} \log m + \log(1/2\epsilon)\right)$. By Corollary 4.4.4 our results show the stronger bound $\tau(\epsilon) \leq 8m n^2 \log \frac{1}{2\sqrt{\pi_*}\epsilon} \leq 8m n^2 \left(\frac{n}{2} \log m + \log(1/2\epsilon)\right)$, exactly the same upper bound obtained by Feder and Mihail [5] by using a modified form of Poincaré.

Also, in [22] we showed that $\Phi(A) \geq \frac{\log_2(1/\pi(A))}{2mn}$, which by equation (3.1) implies that $\tau(\epsilon) \leq 4(\log 2)m^2 n^2 + 8m^2 n^2 \log(1/2\epsilon)$, not particularly good. However, $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq \Phi(A)^2/2$, and combining this with the canonical path lower bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ implies that

$$\begin{aligned} 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) &\geq \max \left\{ \frac{1}{8m n^2}, \frac{\log_2^2(1/\pi(A))}{8m^2 n^2} \right\} \\ &= \begin{cases} \log_2^2(1/\pi(A))/8m^2 n^2 & \text{if } \pi(A) \leq 2^{-\sqrt{m}} \\ 1/8m n^2 & \text{if } \pi(A) > 2^{-\sqrt{m}} \end{cases} \end{aligned}$$

and this time the mixing time is $\tau(\epsilon) \leq 8(\log 2)m^{3/2}n^2 + 8mn^2 \log(1/2\epsilon)$ when $2^{-\sqrt{m}} \geq m^{-n}$, an improvement over the canonical paths bound for all simple balanced matroids (as $m \leq \binom{n}{2}$).

We note that the correct bound is still much smaller at $\tau(\epsilon) = O\left(mn \log \frac{n}{\epsilon}\right)$ [12].

It is apparent from the definitions that $\rho_v \geq \rho_e \min\{P(x, y) : x \neq y, P(x, y) \neq 0\}$, since if \vec{e} is the worst directed edge then v can be taken as one of the endpoints. In the following example we show how a common enumeration process sometimes leads to this lower bound being achieved.

Example 4.4.6. One of the first Markov chains analyzed via the canonical path method was a Markov chain of Broder [1] for approximating the permanent of a dense matrix, or equivalently counting perfect matchings in a bipartite graph $G = K_{n,n}$ of minimum degree $n/2$. Jerrum and Sinclair [11] constructed canonical paths and used a clever enumeration process to show that $\rho_e \leq 12n^6$. We give only a rough sketch of where the ρ_v computation differs from the ρ_e computation, and refer the reader to [11] for further details.

List all cycles in G , assign them an ordering, and to each cycle C fix a “starting vertex” v_c . As done by Jerrum and Sinclair, to each pair of perfect matchings the path γ_{IF} is given by considering the symmetric difference $I \oplus F$, and then unwinding the resulting cycles in the cycle ordering just given. Let M be a fixed vertex (a perfect or near-perfect matching) and suppose γ_{IF} is a canonical path joining two perfect matchings and passing through M . If M is near-perfect then let e_{IM} denote the edge of I incident with the starting vertex of the cycle being unwound when M was reached, and e_{FM} be the edge of F that will be added at the next step of the unwinding. The encoding we use is

$$\sigma_M(I, F) = \begin{cases} I \oplus F \oplus M & \text{if } M \text{ is perfect;} \\ I \oplus F \oplus M - e_{IM} & \text{if } M + e_{IM} \text{ is perfect;} \\ I \oplus F \oplus M - e_{IM} - e_{FM} & \text{otherwise} \end{cases}$$

Checking a few cases, as in [11], shows that $\sigma_M(I, F)$ is a perfect or near-perfect matching.

To see that this is injective we define the decoding process. If $\sigma_M(I, F) \oplus M$ is a sequence of cycles then $I \oplus F = \sigma_M(I, F) \oplus M$. If $\sigma_M(I, F) \oplus M$ is a sequence of cycles plus a path containing the remaining vertices then let e be the edge joining the endpoints of the path, and $I \oplus F = \sigma_M(I, F) \oplus M + e$. The only remaining case is if $\sigma_M(I, F) \oplus M$ is a sequence of cycles, plus two paths; each path will have one end in M and one end not in M , join the paths to create a cycle by connecting the ends in M to those not in M , and this gives $I \oplus F$.

Given $I \oplus F$ then the order in which the unwinding occurred follows from the cycle ordering. The matching I contains all edges removed before M was reached, plus all edges in M for the remaining cycles. The matching F is the converse.

It follows that the enumeration given above measures all paths connecting two perfect matchings and passing through M , not just those through an edge as in [11]. The computation of Jerrum and Sinclair then holds for ρ_v as well, but without the need for the $P(M, M')$ term in the denominator. That is, $\rho_v \leq 12n^6 \max P(M, M') = 6n^4$ and therefore $\lambda \geq 1/288n^{10}$, an improvement over the $\lambda \geq 1/1152n^{12}$ of [11].

The $O(n^2)$ improvement over the results of [11] is nice and is the maximum possible with our path bound, as $\rho_v \geq \rho_e \min\{P(x, y) : x \neq y, P(x, y) \neq 0\}$. However, this is not as large as the $O(n^5)$ improvement possible by use of a Poincaré bound $\lambda \geq 1/\rho_e \ell$, as in [27, 4]. In the following

section a comparison theorem is proven in terms of maximum path length via Blocking conductance, and hence will match the Poincaré bound, at least up to order of magnitude. Unfortunately, we have been unable to derive a similar theorem for the Evolving sets bounds.

The reason our ρ_v based bound is sometimes as good as using path length, sometimes not, and occasionally better, can be understood better by the following corollary.

Corollary 4.4.7. *Let ρ_e and ρ_v be as in Corollary 4.4.4, and also let*

$$\rho_v^{ave} = \sum_{v \in V} \pi(v) \left[\frac{1}{\pi(v)} \sum_{\gamma_{xy} \ni v} \pi(x) \pi(y) \right]$$

be the average vertex congestion over the entire space V . Then

$$\tau(\epsilon) \leq 2 \left(\frac{\rho_v(1 - \|\pi\|_2^2)}{\rho_v^{ave}} \right) \rho_e \ell_{ave} \log \frac{1}{2\epsilon\sqrt{\pi_*}}$$

where $\ell_{ave} = \frac{\sum_{x \neq y \in V} \pi(x) \pi(y) |\gamma_{xy}|}{\sum_{x \neq y \in V} \pi(x) \pi(y)}$ is the average length of the canonical paths.

Proof. Changing the order of summation gives

$$\begin{aligned} \rho_v^{ave} &= \sum_{x \neq y \in V} \sum_{v \in \gamma_{xy}} \pi(x) \pi(y) = \sum_{x \neq y \in V} \pi(x) \pi(y) (|\gamma_{xy}| + 1) \\ &= (\ell_{ave} + 1) \sum_{x \neq y \in V} \pi(x) \pi(y) = (\ell_{ave} + 1)(1 - \|\pi\|_2^2). \end{aligned}$$

Finish by multiplying the upper bound on $\tau(\epsilon)$ in Corollary 4.4.4 by $1 \leq \frac{2\ell_{ave}(1 - \|\pi\|_2^2)}{\rho_v^{ave}}$. □

A similar result for the Comparison theorem also holds, but with the average path length

$$\ell_{ave} = \frac{\sum_{(x,y) \in E'} \pi'(x) P'(x,y) |\gamma_{xy}|}{\sum_{(x,y) \in E'} \pi'(x) P'(x,y)}.$$

This shows that when the canonical paths are well distributed among the vertices, as in Example 4.4.5, then Corollary 4.4.3 can be as strong as Poincaré bounds. The Corollary suggests that if the paths are short but concentrated on a few vertices then our results will be poor. On the other hand, if $\|\pi\|_2 \approx 1$, or if the paths are fairly well distributed among the vertices but there are some paths much longer than average then our vertex bounds may be better. The following example satisfies both of these conditions.

Example 4.4.8. Consider the lazy asymmetric walk on a line $P_n = [1 \dots n]$ given by $P(i, i+1) = \frac{c}{2} \leq 1/2$ and $P(i, i-1) = \frac{1-c}{2} \leq 1/2$, with $P(1, 1) = \frac{1}{2} + \frac{1-c}{2}$ and $P(n, n) = \frac{1}{2} + \frac{c}{2}$. The stationary distribution is $\pi(i) = \beta \alpha^i$ where $\alpha = \frac{c}{1-c}$ and $\beta^{-1} = \sum_{i=1}^n \alpha^i = \frac{c}{1-2c} \left(1 - \left(\frac{c}{1-c} \right)^n \right)$.

In particular, when $c < 1/2$ then $1 - \|\pi\|_2^2 \xrightarrow{n \rightarrow \infty} 2c$, and if $c \approx 0$ then Corollary 4.4.7 suggests that using ρ_v should be much better than using path length ℓ . There is an obvious choice of canonical paths given by $\gamma_{ij} = i, i+1, \dots, j$ when $i < j$, and vice-versa when $i > j$. It is easily verified that $\ell = n-1$, $\rho_e \rightarrow 2$ and $\rho_v \rightarrow 1$ when $c \rightarrow 0^+$ and $n \rightarrow \infty$. Then $\lambda \geq 1/2\rho_e^2 = 1/8$ and $\lambda \geq 1/4\rho_v\rho_e = 1/8$ give the same bounds, whereas a Poincaré inequality implies only that $\lambda \geq 1/\rho_e\ell = 1/2(n-1)$, far worse. The correct value is $\lambda \rightarrow 1/2$.

4.5 An aside into Blocking Conductance and comparison

One advantage of the Spectral Profile method over that of Evolving sets is that it is fairly straightforward to show that one can compare Spectral Profile of two Markov chains, whereas in the Evolving set case we do not know how to compare f -congestions directly. Because of the similarity between Evolving Set and Blocking Conductance bounds we show here an unpublished comparison method for Blocking Conductance, in the hope that it may help give insight into future work in proving such a result for Evolving Sets as well.

Blocking Conductance only applies to reversible chains and the constant factors are significantly weaker than with our evolving set results, so we skip giving proofs in terms of ρ_e and ρ_v as in the previous section, although such bounds are not hard to show. Instead, we only consider a Poincaré type comparison theorem because this is a situation in which we may genuinely improve on known results. In particular, when studying problems for which total-variation mixing time is faster than L^2 mixing, such as the application of Blocking Conductance to studying $G(n, p)$ given by Fountoulakis and Reed [7], a Blocking Conductance comparison theorem might show stronger bounds than the L^2 -based comparison of Spectral Profile method.

Our results will be shown with an earlier form of Blocking Conductance from our Ph.D. Dissertation [19], since it seems best suited for our purposes here.

Given a finite Markov chain \mathcal{M} with state space V of cardinality $n = |V|$, let $(V, <)$ denote an ordering of the space, say as $1, 2, \dots, n$, let A_i denote the subsets $A_i = [1 \dots i]$, and $A_i^+ = A_{i+1} = [1 \dots i + 1]$. Then the blocking conductance theorem says:

Theorem 4.5.1. *Given a lazy, reversible, ergodic, finite Markov chain \mathcal{M} there exists some ordering $(V, <)_{\mathcal{M}}$ of the vertices such that if $S_0 = [1 \dots k]$ is the largest A_i such that $\pi([1 \dots k]) \leq 3/4$, and if $S_0 \supset S_1 \supset \dots \supset S_m = [1]$ is any nested sequence of A_i , then the mixing time is bounded by*

$$\tau(\mathcal{M}) \leq K' \sum_{i=1}^m \frac{\pi(S_i)}{\mathbf{Q}(S_{i+1}^+, S_i^c)},$$

where $K' = 1376$ arises from converting between different measures of mixing time.

Since the ordering is not known in advance it is necessary to maximize the right hand side over all possible orderings. We use the notation

$$\tau_{BC}(\mathcal{M}) = K' \max_{\text{orderings } (V, <)} \min_{\substack{\text{nested sequences} \\ S_0 \supset S_1 \supset \dots \supset S_m = [1]}} \sum_{i=1}^m \frac{\pi(S_i)}{\mathbf{Q}(S_{i+1}^+, S_i^c)}, \quad (4.9)$$

that is, the best possible upper bound on mixing time when the ordering is not known. In Remark 4.5.3 some more conventional forms of this relation are discussed.

Our main result of this section is then:

Theorem 4.5.2 (Comparison with Blocking Conductance). *Suppose that \mathcal{M} and \mathcal{M}' are finite Markov chains with the same set of vertices V , the same stationary distribution π , and suppose that to every edge $e' = (x, y) \in E'$ is associated a path $\gamma_{xy} \subset E$. Let*

$$\rho_e = \max_{e \in E} \frac{1}{\mathbf{Q}(e)} \sum_{\gamma_{xy} \ni e} \pi(x) \mathbf{P}'(x, y)$$

be a measure of edge congestion and $\ell = \max_{(x,y) \in E'} |\gamma_{xy}|$ be the length of the longest path.

Then,

$$\tau_{BC}(\mathcal{M}) \leq 4 \rho_e \ell \tau_{BC}(\mathcal{M}')$$

where τ_{BC} is given by equation (4.9).

Proof. Fix some ordering $(V, <)$ of the state space V , as discussed in the Preliminaries. Let $S_0 \supset S_1 \supset \dots \supset S_m = [1]$ be the nested sequence of sets $A_i = [1 \dots i]$ which minimizes

$$\sum_{i=1}^m \frac{\pi(S_i)}{\mathbf{Q}(S_{i+1}^+, S_i^c)},$$

that is the sum in the upper bound for $\tau_{BC}(\mathcal{M})$. Observe that every set S with $S_1^+ \subseteq S \subseteq S_0$ satisfies $\mathbf{Q}(S, S^c) \geq \mathbf{Q}'(S, S^c)/\rho_e \geq \mathbf{Q}'(S_1^+, S_0^c)/\rho_e$.

Let $B_0 = S_0$, and let B_1 be the smallest set $A_i = [1 \dots i] \subset B_0$ such that $\mathbf{Q}(B_1^+, B_0^c) \geq \mathbf{Q}'(S_1^+, S_0^c)/2\rho_e$. Likewise, let B_2 be the smallest $A_i \subset B_1$ such that $\mathbf{Q}(B_2^+, B_1^c) \geq \mathbf{Q}'(S_1^+, S_0^c)/2\rho_e$. Continue until either $B_i \subseteq S_1$, or until $B_{2\ell}$ is defined.

Suppose that sets up to $B_{2\ell}$ were defined and $B_{2\ell} \supset S_1$. Ergodic flow from set B_{i+1} then satisfies the relation $\mathbf{Q}(B_{i+1}, B_i^c) < \mathbf{Q}'(S_1^+, S_0^c)/2\rho_e$, since B_{i+1} was defined to be the smallest initial sequence $[1 \dots k]$ such that $\mathbf{Q}(B_{i+1}^+, B_i^c) \geq \mathbf{Q}'(S_1^+, S_0^c)/2\rho_e$. By definition of ρ_e the ergodic flow $\mathbf{Q}(B_{i+1}, B_i^c)$ therefore contains under half the canonical paths from S_1^+ to S_0^c , and so more than half these paths pass through $B_i \setminus B_{i+1}$. It follows that if $Cut'(S_1^+, S_0^c)$ denotes the edges passing from S_1^+ to S_0^c in \mathcal{M}' , then

$$\begin{aligned} \sum_{(x,y) \in Cut'(S_1^+, S_0^c)} \pi(x) \mathbf{P}'(x, y) (|\gamma_{xy}| - 1) &\geq \sum_{i=0}^{2\ell-1} \sum_{\substack{(x,y) \in Cut'(S_1^+, S_0^c), \\ \gamma_{xy} \cap (B_i \setminus B_{i+1}) \neq \emptyset}} \pi(x) \mathbf{P}'(x, y) \\ &> 2\ell \mathbf{Q}'(S_1^+, S_0^c)/2 \\ &\geq \sum_{(x,y) \in Cut'(S_1^+, S_0^c)} \pi(x) \mathbf{P}'(x, y) |\gamma_{xy}|. \end{aligned}$$

This gives a contradiction, so it follows that $B_{2\ell} \subseteq S_1$.

By construction, if sets up to $B_{2\ell}$ are defined then

$$\sum_{i=0}^{2\ell-1} \frac{\pi(B_i)}{\mathbf{Q}(B_{i+1}^+, B_i^c)} \leq 4 \rho_e \ell \frac{\pi(S_0)}{\mathbf{Q}'(S_1^+, S_0^c)}. \quad (4.10)$$

If fewer B_i were required then the sum is even smaller. Also, it can be assumed that the last B_i is equal to S_1 , because if say $B_{2\ell} \subsetneq S_1$ then increasing the size to $B_{2\ell} = S_1$ can only decrease the sum in the left side of (4.10), which reinforces the inequality.

Repeating this process for the other sets S_k leads to a sequence of sets in \mathcal{M} which is at most $4 \rho_e \ell$ times worse than the S_k sequence in \mathcal{M}' . This holds for any ordering $(V, <)$ and so the theorem follows. \square

Just as with Diaconis and Saloff-Coste's comparison theorem [3], comparison with the trivial chain $\mathbf{P}'(x, y) = \pi(y)$ gives the familiar bound $\tau(\mathcal{M}) = O(\rho_e \ell \log \frac{1}{\pi_*})$. Therefore, at least as far as

order of magnitude is concerned, Blocking Conductance matches the canonical path bounds found in most applications. However, the comparison method may be superior to other methods when the total variation mixing time is smaller than the L^2 mixing time, such as with $G_{n,p}$ [7].

Remark 4.5.3. *It is easier to understand the Blocking conductance theorem in an alternate form. First, a definition. For $A \subset V$ let*

$$\mathcal{B}(A) = \sup_{\alpha \leq \pi(A)} \min_{\substack{B \subset A, \\ \pi(B) \geq \pi(A) - \alpha}} \frac{\alpha \mathbf{Q}(B, A^c)}{\pi(A)^2}.$$

The quantity $\mathbf{Q}(B, A^c)/\pi(B)$ is a lower bound on the probability of leaving A in a single step, conditioned on the initial point being drawn from B . So $\mathcal{B}(A)$ roughly measures the size of a potential vertex bottleneck relative to the size of A , times the probability of stepping over this bottleneck. For example, if $\alpha = \frac{1}{2} \mathbf{Q}(A, A^c)$ then $\mathbf{Q}(B, A^c) \geq \frac{1}{2} \mathbf{Q}(A, A^c)$ when $\pi(B) \geq \pi(A) - \alpha$, and so $\mathcal{B}(A) \geq \frac{1}{4} \Phi(A)^2$. The name Blocking conductance comes from this connection to conductance and the fact that $\pi(B) \geq \pi(A) - \alpha$ and so α measures the size of a subset of A which is “blocked” when computing $\mathcal{B}(A)$.

Given an ordering $(V, <)$ a good sequence of S_i is easily enough constructed from $\mathcal{B}(\cdot)$. Given S_i then let S_{i+1} be such that $\pi(S_i \setminus S_{i+1}) \approx \alpha$. The following is a slight improvement on a result of the author [19].

Theorem 4.5.4. *If \mathcal{M} is a finite, irreducible, reversible, lazy Markov chain then*

$$\tau(\mathcal{M}) \leq \tau_{BC}(\mathcal{M}) \leq \frac{3}{2} K' \int_{\pi_0/2}^{3/4} \frac{dx}{x \phi(x)}$$

where $K' = 1376$ arises from converting between different measures of mixing time and the blocking conductance function ϕ is given, for $x \in [\pi_0/2, 3/4]$, by

$$\phi(x) = \min_{\substack{\pi(A) \in [x, 2x], \\ \pi(A) \leq 3/4}} \mathcal{B}(A).$$

4.6 Spread and the discrete gradients

The quantity we consider here was proposed by Kannan, Lovász and Montenegro [13] in the context of Blocking Conductance, in an alternative form by Morris and Peres [26], and also used by Montenegro [20]. As a bit of motivation, recall that earlier isoperimetric bounds, such as those involving conductance or modified conductance, were shown by explicitly constructing the worst case for $\pi(A_u)$ and then applying Lemma 3.3.1. We now show that useful relations can be derived even when Lemma 3.3.1 appears not to be applicable. We examine only the relation between the Evolving set bounds and the spread. The interested reader can see Montenegro [20] for an examination of the relation between spread, the discrete gradients of Houdré and Tetali [10], and spectral gap.

Two isoperimetric quantities will be used, extending earlier definitions. Recall from definition (4.4) that $\Psi(A, t)$ is the smallest flow from from A into a subset of V of size t , so in particular for a lazy walk $\Psi(A, \pi(A^c)) = \mathbf{Q}(A, A^c)$.

Definition 4.6.1. If $A \subset V$ then the *internal and external spread* are given by

$$\tilde{\psi}^+(A) = \int_0^{\pi(A)} \frac{\Psi(A^c, t)}{\pi(A)^2 \pi(A^c)^2} dt \quad \text{and} \quad \tilde{\psi}^-(A) = \int_0^{1-\pi(A)} \frac{\Psi(A, t)}{\pi(A)^2 \pi(A^c)^2} dt.$$

Quantities $\psi^\pm(A)$ are defined similarly but without $\pi(A^c)^2$ in the denominators.

The spread turns out to fairly closely bound many natural choices of \mathcal{C}_f . This was first observed in [20] where the connection between $\mathcal{C}_{\sqrt{a}}$ and spread were studied for lazy Markov chains in order to relate Blocking Conductance and Evolving Set results.

Theorem 4.6.2. *Given a finite irreducible Markov kernel and $A \subset V$ with $\pi(A) \leq 1/2$ then*

$$\begin{aligned} \tilde{\psi}^+(A)\pi(A^c) + \tilde{\psi}^-(A) \max\left\{\pi(A), \frac{1}{4}\right\} &\geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq \frac{1}{4} \tilde{\psi}^+(A) + \tilde{\psi}^-(A) \pi(A)\pi(A^c) \\ \frac{2\psi^+(A) + \psi^-(A)}{\log(1/\pi(A))} &\geq 1 - \mathcal{C}_{a \log(1/a)}(A) \geq \frac{\psi^+(A)}{\log(1/\pi(A))} + \tilde{\psi}^-(A) \pi(A)\pi(A^c) \\ \text{and } 1 - \mathcal{C}_{a(1-a)}(A) &= 2 \left(\tilde{\psi}^+(A) + \tilde{\psi}^-(A) \right) \pi(A)\pi(A^c). \end{aligned}$$

The bounds for $1 - \mathcal{C}_{a \log(1/a)}(A)$ and $1 - \mathcal{C}_{a(1-a)}(A)$ hold when $\pi(A) > 1/2$ as well.

It follows that $1 - \mathcal{C}_{a \log(1/a)}(A), 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq \frac{1}{2} (1 - \mathcal{C}_{a(1-a)}(A))$.

Observe that $\Psi(A^c, \pi(A) - t) \geq \Psi(A^c) - t = \Psi(A) - t$, and so

$$\tilde{\psi}^+(A) \geq \frac{\int_0^{\Psi(A)} (\Psi(A) - t) dt}{\pi(A)^2 \pi(A^c)^2} = \frac{1}{2} \tilde{\phi}(A)^2.$$

Likewise, $\tilde{\psi}^-(A) \geq \frac{1}{2} \tilde{\phi}(A)^2$. Therefore, at least up to a small multiplicative factor, this supercede's Theorem 3.4.1.

Remark 4.6.3. *As discussed in Remark 4.3.6, for a lazy walk spread incorporates measures of both edge and vertex expansion, whereas (modified) conductance involves only edge expansion. This can be generalized to a non-lazy walk as well. To see this, given set $A \subset V$ let $\beta(A) = \max\{t : \Psi(A^c, \pi(A) - t) \geq \Psi(A^c)/2\}$. Then $\tilde{\psi}^+(A) \geq \frac{\beta(A)\Psi(A)}{2\pi(A)^2\pi(A^c)^2}$ and so*

$$\tilde{\phi} \geq 1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \geq \min_{\pi(A) \leq 1/2} \frac{1}{4} \tilde{\psi}^+(A) \geq \min_{\pi(A) \leq 1/2} \frac{1}{8} \frac{\beta(A)}{\pi(A)\pi(A^c)} \tilde{\phi}(A).$$

The quantity $\beta(A)$ is a notion of vertex expansion that measures how large a set of vertices must be “blocked” so that only half of the flow $\Psi(A^c)$ remains. Therefore $\tilde{\psi}^+(A)$ can be thought of as a product of edge and vertex expansion. Since $\frac{\pi(A)}{2} \geq \beta(A) \geq \frac{\Psi(A^c)}{2}$ then the above lower bound may be the same order as the upper bound, and is at worst 4 times weaker than our modified conductance lower bound of Theorem 3.4.1. However, when there are many boundary vertices and $\beta(A) \gg \Psi(A^c)$ then this can be substantially better.

Proof. The proof involves working with the spread written in a form involving Evolving Sets:

$$\tilde{\psi}^+(A) = \frac{1}{2} \int_{\wp_A}^1 \frac{(\pi(A_u) - \pi(A))^2}{\pi(A)^2 \pi(A^c)^2} du, \quad \tilde{\psi}^-(A) = \frac{1}{2} \int_0^{\wp_A} \frac{(\pi(A_u) - \pi(A))^2}{\pi(A)^2 \pi(A^c)^2} du$$

The first of these was derived in equation (4.6); the $\tilde{\psi}^-(A)$ relation follows in exactly the same way.

The equality for $1 - \mathcal{C}_{a(1-a)}(A)$ follows immediately from the form in Lemma 4.3.3 and the definitions of $\tilde{\psi}^\pm(A)$. Alternatively, as an introduction to the method used in the remainder of the proof, start with the identity

$$\forall x, y \in [0, 1] : y(1-y) = x(1-x) + (1-2x)(y-x) - (y-x)^2.$$

Letting $y = \pi(A_u)$, $x = \pi(A)$ and integrate over $u \in [0, 1]$. Finish by applying the Martingale property that $\int_0^1 (\pi(A_u) - \pi(A)) du = 0$, and dividing by $\pi(A)\pi(A^c)$.

The inequality required to show the lower bound on $1 - \mathcal{C}_{a \log(1/a)}(A)$ is

$$y \log \frac{1}{y} \leq x \log \frac{1}{x} + \left(\log \frac{1}{x} - 1 \right) (y-x) - \frac{(y-x)^2}{2} \left(\delta_{y \leq x} \frac{1}{x} + \delta_{y > x} \frac{\log \frac{1}{x} - (1-x)}{(1-x)^2/2} \right) \quad (4.11)$$

for all $x, y \in [0, 1]$ (proven below). Let $y = \pi(A_u)$, $x = \pi(A)$, recall $\int_0^1 (\pi(A_u) - \pi(A)) du = 0$, and use the form of $\mathcal{C}_{a \log(1/a)}(A)$ derived in Lemma 4.3.3,

$$\begin{aligned} \mathcal{C}_{a \log(1/a)}(A) &= \int_0^1 \frac{\pi(A_u) \log(1/\pi(A_u))}{\pi(A) \log(1/\pi(A))} du \\ &\leq 1 - \frac{\psi^+(A)}{\log(1/\pi(A))} - 2\tilde{\psi}^-(A) \pi(A)\pi(A^c) \left(\frac{1}{\pi(A^c)} - \frac{1}{\log(1/\pi(A))} \right) \end{aligned}$$

To finish apply the inequality $\forall x \in [0, 1] : \frac{2}{1-x} - \frac{2}{\log(1/x)} \geq 1$ with $x = \pi(A)$.

The upper bound on $1 - \mathcal{C}_{a \log(1/a)}(A)$ follows similarly, but with

$$y \log \frac{1}{y} \geq x \log \frac{1}{x} + \left(\log \frac{1}{x} - 1 \right) (y-x) - \frac{(y-x)^2}{2x} (1 + \delta_{y \leq x})$$

for all $x, y \in [0, 1]$.

For the lower bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ use the inequality

$$\sqrt{y(1-y)} \leq \sqrt{x(1-x)} + \frac{1-2x}{2\sqrt{x(1-x)}} (y-x) - \frac{(y-x)^2}{2[x(1-x)]^{3/2}} \left(\delta_{y \leq x} \frac{1}{4} + \delta_{y > x} x(1-x) \right)$$

for $x \in [0, 1/2]$, $y \in [0, 1]$. This relation follows from the inequality $\sqrt{z} \leq 1 + \frac{1}{2}(z-1) - \frac{1}{8}(z-1)^2 \delta_{z \leq 1}$ with $z = \frac{y(1-y)}{x(1-x)}$. A calculation as done for $\mathcal{C}_{a \log(1/a)}(A)$ gives the result.

The upper bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ uses

$$\sqrt{y(1-y)} \geq \sqrt{x(1-x)} + \frac{1-2x}{2\sqrt{x(1-x)}} (y-x) - \frac{(y-x)^2}{2[x(1-x)]^{3/2}} \left(\delta_{y \leq x} (1-x) + \frac{\delta_{y \in (x, 1-x)}}{4} + \delta_{y \geq 1-x} x \right)$$

for $x \in [0, 1/2]$, $y \in [0, 1]$. This relation follows from the inequality $\sqrt{z} \geq 1 + \frac{1}{2}(z - 1) - \frac{1}{2}(z - 1)^2 \delta_{z \leq 1} - \frac{1}{8}(z - 1)^2 \delta_{z \geq 1}$, again with $z = \frac{y(1-y)}{x(1-x)}$. Consolidate a bit via the relation $\frac{\delta_{y \in (x, 1-x)}}{4} + \delta_{y \geq 1-x} x \leq \delta_{y \geq x} \max\{x, 1/4\}$. Finish again as in $\mathcal{C}_{a \log(1/a)}(A)$.

We finish with a proof of equation (4.11). Consider

$$g(x, y) = y \log \frac{1}{y} - x \log \frac{1}{x} - \left(\log \frac{1}{x} - 1 \right) (y - x) + \frac{(y - x)^2}{2} (h_1(x) \delta_{y \leq x} + h_2(x) \delta_{y > x})$$

where $h_1(x) = 1/x$ and $h_2(x) = \frac{\log \frac{1}{x} - (1-x)}{(1-x)^2/2}$. We will show that $g(x, y) \leq 0$ for $x, y \in [0, 1]$ by showing that $g(x, y)$ is increasing for $y \in [0, x]$, while for $y \in (x, 1]$ it decreases and then increases. It is easily verified that the inequality holds at $y = x$, $y \rightarrow x^+$ and $y = 1$, so the result then follows. First, calculate a few derivatives.

$$\begin{aligned} \frac{dg}{dy} &= \log \frac{x}{y} + (y - x) (h_1(x) \delta_{y < x} + h_2(x) \delta_{y > x}) \\ \frac{d^2g}{dy^2} &= -\frac{1}{y} + h_1(x) \delta_{y < x} + h_2(x) \delta_{y > x} \\ \frac{d^3g}{dy^3} &= \frac{1}{y^2} > 0 \end{aligned}$$

Consider $y \in [0, x)$. Since $h_1(x) = 1/x$ then $\frac{d^2g}{dy^2} \Big|_{y < x} \leq 0$ and so $g(x, y)$ is concave in y . But $\frac{dg}{dy} \Big|_{y \rightarrow x^-} = 0$ and so $g(x, y)$ is increasing in y for $y \in [0, x]$, as desired. Now consider $y \in (x, 1]$. The third derivative is positive, so $\frac{d^2g}{dy^2}$ is increasing in y . But $\frac{d^2g}{dy^2} \Big|_{y \rightarrow x^+} \leq 0$ and $\frac{d^2g}{dy^2} \Big|_{y=1} \geq 0$ and so g is initially concave and transitions to convex. Since $\frac{dg}{dy} \Big|_{y \rightarrow x^+} = 0$ then g decreases, then if it transitions to convex then it may increase later. This completes the proof of equation (4.11). \square

The bounds are sharp.

The worst case for the lower bound of $1 - \mathcal{C}_{a \log(1/a)}(A)$ is when the flow leaves a small sliver of A and flows uniformly into A^c , that is $\Psi(A^c, t) = (t - (\pi(A) - \mathbf{Q}(A, A^c)))^+$ when $t \leq \pi(A)$ and $\Psi(A^c, t) = \frac{1-t}{1-\pi(A)} \mathbf{Q}(A, A^c)$ when $t \geq \pi(A)$. For the upper bound this is reversed, with the flow leaving uniformly from A and concentrated in a sliver of A^c . This is also sharp on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$

The upper bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$ is sharp for all set sizes, despite the odd looking $\max\{1/4, x\}$ term. When $\pi(A) \in [1/4, 1/2]$ then look at the walk on K_n discussed in Example 3.7.3, with $\gamma = 1/n$, that is $\alpha = 0$. Then $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = 1$, $\tilde{\psi}^+(A) = 1/2\pi(A^c)$ and $\tilde{\psi}^-(A) = 1/2\pi(A)$ and so the bound becomes $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) \leq 1/2 + \max\{1/8\pi(A), 1/2\}$ which is correct when $\pi(A) \geq 1/4$. When $\pi(A) < 1/4$ then consider a Markov chain with transitions to be defined. Let A be a set of size $x = \pi(A)$, let $B \subset A^c$ with $\pi(B) = \epsilon$, and let the transition kernel satisfy $\forall v \in A : \mathbf{P}(v, B) = 1 - \mathbf{P}(v, A) = \epsilon/2x$ while $\forall v \in B : \mathbf{P}(v, A) = \mathbf{P}(v, B) = 1/2$. Then it is easily computed that $\tilde{\psi}^+(A) = \frac{\epsilon}{4z(1-z)^2}$, $\tilde{\psi}^-(A) = \frac{\epsilon^2}{4x^2(1-x)^2}$, and $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = \frac{1}{2} + \frac{\epsilon}{2x} - \frac{1}{2} \sqrt{(1 + \epsilon/x)(1 - \epsilon/(1-x))}$. As $\epsilon \rightarrow 0^+$ then $(1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) - \tilde{\psi}^+(A)(1-x))/\tilde{\psi}^-(A) \rightarrow 1/4$ for all $z \in [0, 1/2]$, and so $\tilde{\psi}^-(A)$ must always be multiplied by at least $1/4$ in the upper bound.

Another instance of sharpness is the lazy walk on the line of even length. It is easily checked that $\tilde{\psi}^+(A), \tilde{\psi}^-(A) \geq 1/8n^2 \pi(A)^2 \pi(A^c)^2$, with equality when A is an initial interval of the line. Then the lower bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}$ is achieved at $\pi(A) = 1/2$, with $1 - \mathcal{C}_{\sqrt{a(1-a)}} \geq 1/n^2$. The correct bound is $1 - \mathcal{C}_{\sqrt{a(1-a)}} = \frac{1}{2} \left(1 - \sqrt{1 - 4/n^2}\right) \xrightarrow{n \rightarrow \infty} 1/n^2$, and our bound was correct.

As mentioned earlier, the spread incorporates measures of both edge and vertex expansion, whereas (modified) conductance involves only edge expansion. Hence an improvement will be likely when vertex expansion is much larger than edge expansion. The most extreme example of this is a walk on the complete graph.

Example 4.6.4. Consider the lazy Markov chain on the complete graph K_n given by choosing a vertex uniformly at random, and moving there with probability $1/2$.

Conductance based bounds tend to be decent when considering the L^2 distance, but may be poor for other distances. In contrast, the χ^2 bound is fine. Likewise, in this case the Cheeger inequality of Theorem 3.4.1 is fine for the bound on $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A)$. In contrast, $1 - \mathcal{C}_{a(1-a)}(A) = 1 - \mathcal{C}_{a \log(1/a)}(A) = 1/2$ but the Cheeger inequalities of Theorem 3.4.1 shows only $1 - \mathcal{C}_{a(1-a)}(A) \geq \pi(A)\pi(A^c)$ and $(1 - \pi(A))^2/2 \log(1/\pi(A))$, both of which go to 0 as $\pi(A) \rightarrow 0^+$.

We now use the spread. It is clear that $\Psi(A^c, t) = \frac{\pi(A^c)}{2}t$ if $t \leq \pi(A)$, while $\Psi(A, t) = \frac{\pi(A)}{2}t$ if $t \leq \pi(A^c)$, and so $\tilde{\psi}^+(A) = \frac{1}{4\pi(A^c)}$ while $\tilde{\psi}^-(A) = \frac{1}{4\pi(A)}$. The lower bounds are now within a factor of two, with $1 - \mathcal{C}_{a(1-a)}(A) = 1/2$ and $1 - \mathcal{C}_{a \log(1/a)}(A) \geq \frac{1}{4}$.

Compare this to the lazy random walk on the cycle of odd length (see Example 3.7.4). In this case $\mathbf{Q}(A, A^c) = 1/n$ and $\beta(A) = 1/n$, so edge and vertex expansion are of similar orders. Therefore it is not surprising that the lower bounds on $\mathcal{C}_{a(1-a)}$ and $\mathcal{C}_{a \log(1/a)}$ given by Theorem 3.4.1 will be the correct order.

Another case where vertex expansion is high is a walk on a product space. See [20] for a proof that the lazy walk on a Boolean cube 2^d has $\tilde{\psi}^+(A) = \Omega\left(\frac{1}{d \log d}\right)$, and so $1 - \mathcal{C}_{\sqrt{a(1-a)}}(A) = \Omega\left(\frac{1}{d \log d}\right)$, a substantial improvement on the Cheeger inequality bound of $\Omega(1/d^2)$, although when it comes to mixing time this is still not as good as what can be shown by log-Sobolev or spectral methods.

Bibliography

- [1] A.Z. Broder. How hard is it to marry at random? (on the approximation of the permanent). *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC 1986)*, pages 50–58, 1986.
- [2] P. Diaconis and J. Fill. Strong stationary times via a new form of duality. *The Annals of Probability*, 18(4):1483–1522, 1990.
- [3] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.
- [4] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of markov chains. *The Annals of Applied Probability*, 1:36–61, 1991.
- [5] T. Feder and M. Mihail. Balanced matroids. *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC 1992)*, pages 26–38, 1992.
- [6] J. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The Annals of Applied Probability*, 1(1):62–87, 1991.
- [7] N. Fountoulakis and B. Reed. The evolution of the mixing rate. preprint available online at arXiv:math/0701474.
- [8] N. Fountoulakis and B. Reed. Faster mixing and small bottlenecks. *Probability Theory and Related Fields*, 137:475–486, 2007.
- [9] S. Goel, R. Montenegro, and P. Tetali. Mixing time bounds via the spectral profile. *Electronic Journal of Probability*, 11:1–26, 2006.
- [10] C. Houdré and P. Tetali. Isoperimetric invariants for product markov chains and graph products. *Combinatorica*, 24(3):359–388, 2004.
- [11] M. Jerrum and A. Sinclair. Conductance and the rapid mixing property for markov chains: the approximation of the permanent resolved. *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, pages 235–243, 1988.
- [12] M. Jerrum and J-B. Son. Spectral gap and log-sobolev constant for balanced matroids. *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2002)*, pages 721–729, 2002.

- [13] R. Kannan, L. Lovász, and R. Montenegro. Blocking conductance and mixing in random walks. *Combinatorics, Probability and Computing*, 15(4):541–570, 2006.
- [14] L. Lovász and R. Kannan. Faster mixing via average conductance. *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC 1999)*, pages 282–287, 1999.
- [15] L. Lovász and M. Simonovits. Random walks in a convex body and an improved volume algorithm. *Random Structures and Algorithms*, 4:359–412, 1993.
- [16] M. Mihail. Conductance and convergence of markov chains—a combinatorial treatment of expanders. *30th Annual Symposium on Foundations of Computer Science*, pages 526–531, 1989.
- [17] R. Montenegro. The simple random walk and max-degree walk on a directed graph. preprint at arXiv:math/0609303.
- [18] R. Montenegro. Two conductance theorems, two canonical path theorems, and two walks on directed cayley graphs. preprint at arXiv:math/0611585.
- [19] R. Montenegro. *Faster Mixing by Isoperimetric Inequalities*. Ph.d. thesis, Department of Mathematics, Yale University, 2002. PDF available on the author’s web page at <http://www.ravimontenegro.com/research/thesis.pdf>.
- [20] R. Montenegro. Vertex and edge expansion properties for rapid mixing. *Random Structures and Algorithms*, 26(1–2):52–68, 2005.
- [21] R. Montenegro. Sharp edge, vertex, and mixed cheeger inequalities for finite markov kernels. *Electronic Communications in Probability*, 12:377–389, 2007.
- [22] R. Montenegro and J-B. Son. Edge isoperimetry and rapid mixing on matroids and geometric markov chains. *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 704–711, 2001.
- [23] R. Montenegro and P. Tetali. *Mathematical Aspects of Mixing Times in Markov Chains*, volume 1:3 of *Foundations and Trends in Theoretical Computer Science*. NOW Publishers, Boston-Delft, June 2006.
- [24] B. Morris. The mixing time for simple exclusion. *Annals of Applied Probability*, 16(2):615–635, 2006.
- [25] B. Morris. The mixing time of the thorp shuffle. *SIAM Journal on Computing (SICOMP)*, 2006. to appear.
- [26] B. Morris and Y. Peres. Evolving sets, mixing and heat kernel bounds. *Probability Theory and Related Fields*, 133(2):245–266, 2005.
- [27] A. Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992.