CS860 Winter 2022 Project Report
Topic: Flow Techniques for Graph Expansion

Kam Chuen (Alex) Tung

April 16, 2022

Contents

1 Introduction

Flow techniques have been proven useful in approximating and upper bounding expansion parameters of graphs, notably the edge expansion $\phi'(G)$. On the one hand, there is one fundamental and well-known connection between flows and cuts, via the max flow/min cut theorem for s-t flows. On the other hand, this connection alone does not seem to justify the success of flow-based algorithms and arguments, and surely there are more reasons that flows appear everywhere.

In this project, we investigate instances that flow techniques were employed to approximate and upper bound certain graph expansion parameters: edge expansion $\phi'(G)$, vertex expansion $\psi(G)$, second eigenvalue $\lambda'_2(G)$ of the unnormalized Laplacian, and second reweighted eigenvalue $\lambda^*_2(G)$. The hope is to identify recurring themes that cement the role of flows in these instances of work, and to adapt the results for one parameter to say something about another parameter (e.g. from $\phi'(G)$ to $\psi(G)$, from $\lambda'_2(G)$ to $\lambda^*_2(G)$).

The project consists of three main sections, which can be read independently:

- In Section 2, we prove an inequality relating the vertex expansion $\psi(G)$ of a graph G and the second smallest eigenvalue $\lambda'_2(G)$ of its unnormalized Laplacian. The proof, due to Alon [Alo86], considers an s-t flow network with appropriate edge capacities and uses max flow/min cut duality.
- In Section 3, we introduce the cut-matching game for approximating edge expansion $\phi'(G)$. The principle is that multicommodity flows can be used to certify the expansion of a graph. The new idea of cut-matching game is that the flow demand graph is constructed iteratively, by taking unions of perfect matchings. After describing the cut-matching game and showing how to use it to approximate $\phi'(G)$, we show how to adapt the proof for approximating vertex expansion $\psi(G)$. The adaptation is simple and not unknown to the community (see e.g. [CS21]), but the strongest form of the result has not been explicitly stated and proved.
- In Section 4, we see how flow-based techniques can be employed to upper-bound $\lambda'_2(G)$ of the unnormalized Laplacian for special classes of graphs. In order of generality, these classes are: planar graphs, graphs of bounded genus g, and graphs that excludes K_h as a minor. The bounds on $\lambda'_2(G)$ are due to Biswal, Lee, and Rao [BLR10]. The main idea is to relate the eigenvalues to a metric spreading quantity, then taking dual to obtain a multicommodity flow problem where the objective is to minimize some measure of congestion. Then, we show how to adapt their approach to upper bound the reweighted second eigenvalue $\lambda^*_2(G)$. We also include a proof of the planar separation theorem via reweighted eigenvalues. These results are new to the best of our knowledge.

Notations. Given a graph $G = (V, E)$. Unless otherwise specified, we use $n := |V|$ for the number of vertices and $m := |E|$ for the number of edges.

For $S, T \subseteq V$, let $N(S) := \{v \in V \setminus S : (u, v) \in E \text{ for some } u \in S\}$ be the (outer) neighbourhood of S, $E(S, T) := \{(u, v) \in E : u \in S, v \in T\}$ be the set of edges between S and T, and $vol(S) := \sum_{u \in S} deg(S)$ be the volume (total degree) of S.

We use $\phi(G)$ for (edge) conductance, $\phi'(G)$ for edge expansion, and $\psi(G)$ for vertex expansion of a graph G.[1]Writing $A = A(G)$ to be the adjacency matrix of G and $D = D(G)$ to be the (diagonal) degree matrix of G, we let $0 \le \lambda_1(G) \le \lambda_2(G) \le \cdots \le \lambda_n(G)$ be the eigenvalues of the normalized Laplacian $L = I - D^{-1/2} A D^{-1/2}$ and $0 \le \lambda'_1(G) \le \lambda'_2(G) \le \cdots \le \lambda'_n(G)$ be the eigenvalues of the unnormalized Laplacian $L' = D - A$. When the context is clear, we drop the G and write ϕ, λ'_2, etc.

[1] Recall that for $S \subseteq V$, $\phi(S) := |E(S, S^c)|/vol(S)$, $\phi'(S) := |E(S, S^c)|/|S|$, and $\psi(S) := |N(S)|/|S|$. In the course notes Φ is used instead of ϕ' for edge expansion, but I prefer using $'$ for the "unnormalized" quantities – hence ϕ' and λ'_k. (We never take derivatives, so there is no ambiguity.)

$X \lesssim Y$ means $X = O(Y)$ and $X \gtrsim Y$ means $X = \Omega(Y)$.

Flow preliminaries. There are two ways to look at a flow problem: one of maximizing the amount of flow sent, and one of minimizing congestion.

For an s-t flow problem, the first viewpoint corresponds to sending the maximum amount of flow possible from s to t, while respecting the edge (resp. vertex) capacity constraints. That is, the total amount of flow passing through each edge (resp. vertex) does not exceed the amount prescribed for that edge (resp. vertex). Unless otherwise specified (e.g. in Section 3.3), edges have capacity constraints while vertices do not.

The second viewpoint corresponds to sending a prescribed amount of flow from s to t as given by the demand D, while minimizing the maximum edge (resp. vertex) congestion, or some other measures of congestion. In the context of this report, we assume that edges (resp. vertices) have unit capacity. Then, the congestion of an edge (resp. a vertex) is just the total amount of flow passing through it. Again, by default we assume that we are dealing with edge congestion instead of vertex congestion. (Vertex congestion features in Sections 3.3 and 4.4.)

For an s-t flow problem (with the first viewpoint), one of the most important results is the max flow/min cut theorem, which states that the maximum amount of s-t flow that can be sent is equal to the minimum total cost of edges to be cut so that s and t are disconnected. When we are looking at edge (resp. vertex) cuts, the cost of the cut is the sum of capacities of the edges (resp. vertices) that are cut.

The s-t flow problem is sometimes called a single-commodity flow problem, to contrast with its generalization – multicommodity flow problem (MFP). As the name suggests, an MFP may have many commodities: source-sink pairs (s_i, t_i) with demands $D(s_i, t_i)$. The task is to *simultaneously* send $\mathcal{F} \cdot D(s_i, t_i)$ units of flow from s_i to t_i, while respecting the edge/vertex capacities. The maximum $\mathcal{F}$ such that this is possible is defined as the max flow of the MFP. The corresponding congestion minimization problem is to simultaneously send $D(s_i, t_i)$ units of flow from s_i to t_i, while minimizing some measure of congestion.

Associated with an MFP is the *demand graph* $H = H(D)$. It is a weighted, undirected graph. For each commodity (s_i, t_i), there is an edge connecting s_i and t_i with weight $D(s_i, t_i)$. We say that H can be *embedded* in G with edge (resp. vertex) congestion $c > 0$ if there is a flow solution that satisfies all the demand pairs, and its maximum edge (resp. vertex) congestion is at most c.

A *uniform MFP* is when the demand graph is the complete graph K_n with $D(u, v) = 1$ for all $u \neq v \in V$. The max flow problem is then to maximize $\mathcal{F}$, such that it is possible to simultaneously send $\mathcal{F}$ units of flow between each pair of vertices, while respecting the edge capacity constraints.

In Section 3, we will use multicommodity flows to approximate $\phi'(G)$. The underlying principle is that the demand graph H of the MFP serves as a certificate of expansion. We state it as a lemma:

Lemma 1.1 (Demand graph certifies $\phi'(G)$)**.** *Let G be a graph, and let H be the demand graph of an MFP which can be embedded in G with edge congestion c. Then, $\phi'(G) \geq \phi'(H)/c$.*

Proof. Let $S \subseteq V$ such that $|S| \leq n/2$. We want to show that $\phi'(S) \geq \phi'(H)/c$.

From set S, H demands that we send at least $\phi'(H) \cdot |S|$ units of flow out of S. Each outgoing flow path must go through $E(S, S^c)$, so the total congestion over all edges in $E(S, S^c)$ is at least $\phi'(H) \cdot |S|$. The feasibility of the flow problem when edge capacity is c implies that

$$|E(S, S^c)| \cdot c \geq \phi'(H) \cdot |S|.$$

Rearranging and minimizing over S, we are done. □

2 Classical vertex expansion inequalities

Let us warm up with a classical result connecting $\psi(G)$ and $\lambda_2'(G)$.

Theorem 2.1 ([Tan84], [AM85], [Alo86])**.** *Let $G = (V, E)$ be a graph with maximum degree d. Then,*

$$\psi(G) \geq \frac{2\lambda_2'(G)}{d + 2\lambda_2'(G)} \quad \text{and} \quad \lambda_2'(G) \geq \frac{\psi(G)^2}{4 + 2\psi(G)^2}.$$

The first inequality is the "easy" direction proved by Tanner [Tan84] and Alon and Milman [AM85], and the second inequality is the "hard" direction proved by Alon [Alo86]. The proof strategy for the easy direction is to construct an appropriate test vector given a vertex cut, and we will not delve into the detail. It is the hard direction that is of interest here. We shall see how to round the second smallest eigenvector f to a vertex cut using flows. The proof presented here follows that of [Alo86].

For presentation, we outline the steps here and defer some of the proofs to Appendix A.

Step 1: Truncating f

First, we need to truncate f so that $|\operatorname{supp}(f)|$ is at most $n/2$. This is to ensure that the cut produced, which will be a subset of $\operatorname{supp}(f)$, will have size at most $n/2$. A standard strategy is to take either the positive part or the negative part of f and show that the Rayleigh quotient remains small. Write

$$R'(f) := \frac{\sum_{(u,v)\in E}(f(u) - f(v))^2}{\sum_{u\in V} f(u)^2}.$$

Lemma 2.2 (Truncation)**.** *Given a graph $G = (V, E)$, if f is an eigenvector of the unnormalized Laplacian $L'(G)$ with eigenvalue $\lambda_2'(G)$, then letting $f_+ := \max(f, 0)$ and $f_- := \max(-f, 0)$, we have $\max(R'(f_+), R'(f_-)) \leq R'(f) = \lambda_2'(G)$.*

Step 2: Constructing flow network

Now we have a vector $g = f_+$ or f_- such that $g \geq 0$, $|\operatorname{supp}(g)| \leq n/2$ and $R'(g) \leq \lambda_2'$. We construct the following s-t flow problem:

- On the s side, let $A \cong \operatorname{supp}(g)$. For each $u \in A$, connect s and u via an edge with capacity $1 + \psi$, where $\psi = \psi(G)$ is the vertex expansion of G.
- On the t side, let $B \cong V$. For each $v \in B$, connect v and t via a unit-capacity edge.
- If $u \in A$ and $v \in B$, connect u and v via a unit-capacity edge iff $u = v$ or $(u, v) \in E$.

As an example, Figure 1 shows a graph and the corresponding flow network for given $\operatorname{supp}(g)$.

The largest possible value of the flow problem is $(1 + \psi)|A|$, because this is the total capacity of edges going out of s. Indeed we shall prove that:

Lemma 2.3 (Full flow)**.** *The value of the flow problem defined above is exactly $(1 + \psi)|A|$.*

Qualitatively we expect this to be true, because the amount of flow we wish to send depends on ψ, and the larger ψ is, the more flow we should be able to send across.

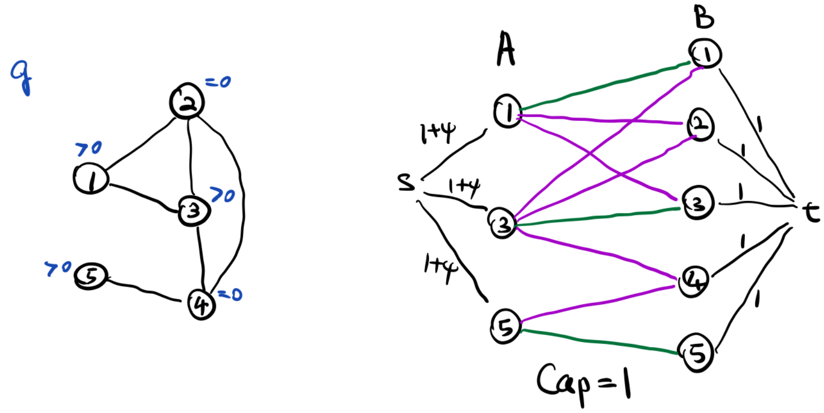

Figure 1: The graph G is shown on the left. $\mathrm{supp}(g) = \{1, 3, 5\}$, so $A = \{1, 3, 5\}$ on the left side of the flow network. The green edges are between $u \in A$ and $u \in B$, and the purple edges are between $u \in A$ and $v \in B$ for all $(u, v) \in E$.

Step 3: Relating Rayleigh quotient and flow solution

Our goal has been to lower bound $R'(g)$, and to this end we shall use the flow network. Fix a solution to the flow problem with value $(1 + \psi)|A|$. For $(u, v) \in E$, if $u \in A$ and $v \in B$, let $h(u, v)$ be the amount of flow sent from u to v; otherwise let $h(u, v) = 0$. Note that $0 \le h(u, v) \le 1$, and also

$$\forall u \in A,\ \sum_{v \in B} h(u, v) = 1 + \psi \quad \text{and} \quad \forall v \in B,\ \sum_{u \in A} h(u, v) \le 1.$$

We need the following two lemmas. The desire to apply Cauchy-Schwarz might be the motivation for introducing the expressions on the LHS's.

Lemma 2.4 (Flow inequality I)**.** $\sum_{(u,v) \in E} h(u, v)^2 (g(u) + g(v))^2 \le (4 + 2\psi^2) \sum_{u \in V} g(u)^2$.

Lemma 2.5 (Flow inequality II)**.** $\sum_{(u,v) \in E} h(u, v)(g(u)^2 - g(v)^2) \ge \psi \sum_{u \in V} g(u)^2$.

Now we are ready to prove the hard direction of Theorem 2.1. Piecing together everything we know,

$$\begin{aligned}
\lambda_2' &\ge R'(g) \qquad \text{(Lemma 2.2)} \\
&= \frac{\sum_{(u,v) \in E} (g(u) - g(v))^2}{\sum_{u \in V} g(u)^2} \\
&= \frac{\sum_{(u,v) \in E} (g(u) - g(v))^2 \cdot \sum_{(u,v) \in E} h(u, v)^2 (g(u) + g(v))^2}{\sum_{u \in V} g(u)^2 \cdot \sum_{(u,v) \in E} h(u, v)^2 (g(u) + g(v))^2} \\
&\ge \frac{\left[\sum_{(u,v) \in E} h(u, v)(g(u) - g(v))(g(u) + g(v))\right]^2}{\sum_{u \in V} g(u)^2 \cdot \sum_{(u,v) \in E} h(u, v)^2 (g(u) + g(v))^2} \qquad \text{(Cauchy-Schwarz)} \\
&\ge \frac{\left[\psi \sum_{u \in V} g(u)^2\right]^2}{\sum_{u \in V} g(u)^2 \cdot (4 + 2\psi^2) \sum_{u \in V} g(u)^2} \qquad \text{(Lemmas 2.4 and 2.5)} \\
&= \frac{\psi^2}{4 + 2\psi^2}.
\end{aligned}$$

3 Cut-matching game

3.1 The cut-matching game framework

The cut-matching game is first devised by Khandekar, Rao, and Vazirani [KRV09] to obtain an $O(\log^2 n)$-approximation algorithm of $\phi'(G)$, which is subsequently improved to $O(\log n)$ by Orecchia et al. [Ore+08]. The motivation for its conception is to obtain a *fast* algorithm for approximating $\phi'(G)$. As we have seen in the preliminaries section, the demand graph H of a multicommodity flow problem (MFP) can be used to certify expansion of G. In particular, if each edge of G has unit capacity, and the MFP has congestion c, then $\phi'(G) \geq \phi'(H)/c$.

In the cut-matching game, the goal is to build a demand graph that has good expansion. The demand graph is built by taking the union of perfect matchings. Note that the edges of the matchings do not have to belong in E. The matchings are constructed via a two-player game that involves a cut player and a matching player. The game consists of T rounds, where T will be specified later. In each round $i \leq T$, the cut player picks a perfect bipartition $X_i \sqcup Y_i$ of the vertex set V (so that $|X_i| = |Y_i| = n/2$) and sends it to the matching player. Then, the matching player chooses a perfect matching M_i between X_i and Y_i.

On the one hand, the task of the cut player is to ensure that, with decent probability, the union of the matchings

$$H_T := \sqcup_{i \leq T} M_i$$

has good edge expansion. (We count repeated edges, so H_T should be regarded as a multigraph.) To this end, he must choose his cuts carefully, so that no matter what the matching player does, H_T will have good expansion. In other words, he adopts a worst-case analysis that assumes adversarial behaviour of the matching player.

On the other hand, the task of the matching player is to return a matching so that the associated flow problem has low congestion: the congestion c_i in each round i should be at most $1/\phi'(G)$. We can reduce it to checking whether or not the max flow of an s-t flow problem with prescribed edge capacities attains a target value. In the end, the maximum congestion $c_{\max} := \max_{i \leq T} c_i$ across all rounds will be used to approximate $\phi'(G)$.

3.2 $O(\log^2 n)$ and $O(\log n)$ approximations of $\phi'(G)$

The bulk of the work in [KRV09] and [Ore+08] lies in devising a good cut player strategy and proving that it works. Indeed, the difference between the approximation ratios of $O(\log^2 n)$ and $O(\log n)$ is solely due to different cut player strategies. We will however only provide a brief description of the cut player strategy, treating the expansion guarantee of H_T as a black-box, and instead focus on the matching player and on how to connect everything to yield the desired approximation of $\phi'(G)$.

Cut player. The role of the cut player can be summarized in the following result:

Theorem 3.1 ([Ore+08, Theorem 3.2]). *In $T = O(\log^2 n)$ rounds of cut-matching game, where the cut player follows strategy C_{nat} as detailed in the paper, and the matching player follows any strategy, returning matchings M_i of size $n/2$ across the cut, the union of the matchings*

$$H_T := \sqcup_{i \leq T} M_i$$

satisfies $\phi'(H_T) \geq \Omega(\log n)$ with constant probability.

In [KRV09], the number of rounds T is the same ($T = O(\log^2 n)$), but the expansion guarantee is weaker, only that $\phi'(H_T) \geq \Omega(1)$. In their paper, the cut player generates the cuts as follows: in round $i \leq T$, start with a random vector $f : V \to \mathbb{R}$ and do random walk $M_1, M_2, \ldots, M_{i-1}$ successively, where the walk M_j has the effect of replacing $f(u)$ and $f(v)$ by their average, for all matching edges $(u, v) \in M_j$. In the end, take X_i to be the $n/2$ vertices with the smallest f value and Y_i to be the rest. The random walk can be implemented in $O(nT^2)$ time, and since $T = O(\log^2 n)$ the overall runtime for the cut player is $\tilde{O}(n)$.

In [Ore+08], a different cut player strategy improves the expansion guarantee of H_T to $\phi'(H_T) \geq \Omega(\log n)$. Again, in round $i \leq T$, the cut player starts with a random vector $f : V \to \mathbb{R}$. Let

$$N_k := \frac{T}{T+1} I + \frac{1}{T+1} M_k \quad \text{and} \quad R_i := N_i N_{i-1} \cdots N_2 N_1 N_2 \cdots N_{i-1} N_i.$$

The cut player applies R_i^T to the vector f, then takes X_i to be the $n/2$ vertices with the smallest f value and Y_i to be the rest.

The random walk R_i is closely related to the natural lazy random walk

$$W_i := \frac{I}{2} + \frac{1}{2i} \sum_{j \leq i} M_j,$$

and the reason for introducing R_i is for ease of analysis. (In their paper, this strategy is denoted $\mathcal{C}_{\text{nat}}$. They propose another strategy, denoted $\mathcal{C}_{\text{exp}}$, which uses matrix exponential to define the random walk. Refer to Section 3 of the paper for more details.) Again, the runtime for the cut player can be shown to be $\tilde{O}(n)$.

Matching player. In round $i \leq T$, once he receives the bipartition $V = X_i \sqcup Y_i$, the matching player attempts to find a matching M_i, such that, if M_i is treated as the demand graph of a multicommodity flow problem, then there is a way to send the flows such that the maximum edge congestion is "small". We describe the construction of the s-t flow network and state the congestion upper bound result. The proof is deferred to Appendix B.

Formally, we build an s-t flow network F_i, where:

- There is a unit capacity edge (s, u) for each $u \in X_i$ and (v, t) for each $v \in Y_i$;
- The edges of the graph G are added to the network, each with capacity c, where c is yet to be specified.

Then, the edge congestion c_i is defined to be the smallest c such that it is possible to send exactly $n/2$ units of flow from s to t. Note that $c_i \geq 1$.

Lemma 3.2 (Edge cut and congestion). *In each round i, we can find a set $S_i \subseteq V$ such that $|S_i| \leq n/2$ and $\phi'(S_i) \leq 1/c_i$.*

Let $c_{\max} := \max_{i \in T} c_i$. An immediate corollary is that:

Corollary 3.3. $\phi'(G) \leq 1/c_{max}$.

Relating $\phi'(H_T)$ and $\phi'(G)$. The following lemma relates the edge expansion of the constructed graph H_T and the edge expansion of G. Namely,

Lemma 3.4 (Certifying edge expansion). $\frac{\phi'(H_T)}{\sum_i c_i} \leq \phi'(G)$.

Proof. Since the MFP with demand graph M_i can be embedded in G with congestion c_i, adding up the flow solutions, the MFP with demand graph $H_T := \sqcup_{i\le T} M_i$ can be embedded in G with congestion at most $\sum_i c_i$. The result now follows from Lemma 1.1. $\square$

Piecing all together. Combining Lemma 3.1, Corollary 3.3, and Lemma 3.4, we have

$$\frac{\phi'(H_T)}{T\cdot c_{\max}} \le \frac{\phi'(H_T)}{\sum_i c_i} \le \phi'(G) \le \frac{1}{c_{\max}}.$$

Since $\frac{\phi'(H_T)}{T} \gtrsim \frac{\log n}{\log^2 n} = \frac{1}{\log n}$, it follows that we have an $O(\log n)$-approximation of $\phi'(G)$.

Runtime analysis. As we have seen, the cut player strategy can be implemented in $\tilde{O}(n)$ time. The matching player strategy consists of computing polylogarithmically many s-t flow problems[2], so it takes $\tilde{O}(T_{\text{flow}})$ time. Undirected flows can be approximated up to constant factor in $O(m^{1+o(1)})$ time, due to Kelner et al. [Kel+14]. Therefore, we arrive at an overall runtime of $O(m^{1+o(1)})$ (we absorb the polylogarithmic factor in the $m^{o(1)}$ term).

3.3 $O(\log n)$ approximation of $\psi(G)$

It turns out that, by using the cut player strategy in [Ore+08] and adapting their proof to the vertex setting, we can obtain an $O(\log n)$ approximation of $\psi(G)$, and the approximation algorithm needs only to compute $O(\log^2 n)$ vertex-capacitated s-t flows.

I thank Robert Wang for working out the details together and Yueheng Zhang and Lap Chi Lau for the discussions. I thank also Thatchaphol Saranurak for confirming the correctness of the result.

Cut player. The cut player follows the same strategy as when approximating $\phi'(G)$.

Matching player. The change here is that we consider vertex congestion instead of edge congestion. The s-t flow network F_i is constructed as follows:

- There is a unit capacity edge (s, u) for each $u \in X_i$ and (v, t) for each $v \in Y_i$;
- Each vertex $u \in X_i \cup Y_i$ has capacity c, which is yet to be specified;
- The edges of the graph G are added to the network, with infinite capacity.

Then, the *vertex congestion* c_i is defined to be the smallest c such that it is possible to send exactly $n/2$ units of flow from s to t. Note that $c_i \ge 1$.

The following lemma upper bounds c_i in terms of vertex expansion. Its proof is deferred to Appendix B.

Lemma 3.5 (Vertex cut and congestion)**.** *In each round i, we can find a set $S_i \subseteq V$ such that $|S_i| \le n/2$ and $\psi(S_i) \le 1/(c_i - 1)$.*

Let $c_{\max} := \max_{i\le T} c_i$. As a corollary of Lemma 3.5 we have

Corollary 3.6. $\psi(G) \le 1/(c_{\max} - 1)$.

[2] There is a technical point about computing c_i. Since $1 \le c_i \le n/2$, we can do binary search on c_i to obtain a constant factor approximation of c_i, which suffices for our purpose. Therefore, each iteration the cut player computes $O(\log n)$ s-t flows.

Relating $\phi'(H_T)$ and $\psi(G)$. The following lemma relates the edge expansion of the constructed graph H_T and the *vertex* expansion of $\psi(G)$. Namely,

Lemma 3.7 (Certifying vertex expansion). $\frac{\phi'(H_T)}{\sum_i c_i} \leq \psi(G)$.

The proof can be found in Appendix B.

Piecing all together. Combining Lemma 3.1, Corollary 3.6, and Lemma 3.7, we have

$$\frac{\phi'(H_T)}{T \cdot c_{\max}} \leq \frac{\phi'(H_T)}{\sum_i c_i} \leq \psi(G) \leq \frac{1}{c_{\max} - 1}.$$

Since $\frac{\phi'(H_T)}{T} \geq \frac{1}{\log n}$, this means we have an $O(\log n \cdot \frac{c_{\max}}{c_{\max}-1})$-approximation of $\psi(G)$. If $c_{\max} \geq 2$, we have an $O(\log n)$ approximation. If $c_{\max} < 2$, this means

$$\frac{1}{\log n} \lesssim \frac{\phi'(H_T)}{T \cdot c_{\max}} \leq \psi(G),$$

and since $\psi(G) \leq O(1)$, we still get an $O(\log n)$ approximation of $\psi(G)$.

Runtime analysis. The cut player strategy can be implemented in $\tilde{O}(n)$ time. The matching player strategy consists of computing polylogarithmically many vertex-capacitated s-t flow problems. Since vertex-capacitated s-t flows can be reduced to directed flows, they can be computed in $O(m^{1+o(1)})$ time, for example by using a recent result of Chen et al. [Che+22] about fast directed flow computations. The matching player runtime, and hence the overall runtime of the algorithm, is thus $O(m^{1+o(1)})$. We have finally proved that:

Theorem 3.8 (Cut matching game for $\psi(G)$). *There exists a randomized algorithm with runtime $O(m^{1+o(1)})$ that, given any graph $G = (V, E)$, computes an $O(\log n)$-approximation of $\psi(G)$ with $\Omega(1)$ probability.*

4 Flows for upper bounds

Now, let us switch gears and look at how flow techniques can be used to upper bound graph expansion parameters. This section is longer than the others because we will be presenting a few new results.

4.1 Background

In 1996, Spielman and Teng [ST96] proved that:

Theorem 4.1 ([ST96, Theorem 3.3]). *For any planar graph G with maximum degree d,*

$$\lambda_2'(G) \leq O(d/n).$$

If the graph G is planar and has bounded degree, then their result combined with Cheeger's inequality implies that G has a balanced edge separator (i.e. a bipartition $V = X \sqcup Y$ of the graph where $|X|, |Y| \geq \Omega(n)$) with $O(\sqrt{n})$ edges in between, hence a balanced vertex separator (i.e. a tripartition $V = A \sqcup B \sqcup C$ of the graph where $|A|, |C| \geq \Omega(n)$ and $E(A, C) = \varnothing$) with $O(\sqrt{n})$ cut vertices, recovering the planar separator theorem of Lipton and Tarjan [LT79] in the bounded-degree case. More importantly, their result explains the success of spectral partitioning when the graph is planar.

In the paper, they conjectured that similar eigenvalue bounds hold for graphs with bounded genus g and graphs which are K_h-minor free. The genus of a graph G is the smallest integer $g \geq 0$ such that G can be embedded in an orientable genus g surface (i.e. torus with g holes) without crossing edges. A graph G contains another graph H as a minor if we can obtain H from G by (1) contracting edges, (2) deleting edges, and (3) deleting vertices, otherwise we say that G is H-minor free. K_h is the complete graph on h vertices.

Kelner [Kel06] first proved the conjecture for graphs with bounded degree and bounded genus.

Theorem 4.2 ([Kel06, Theorem 2.3])**.** *For any graph G with genus $g \geq 1$ and maximum degree d,*

$$\lambda_2'(G) \leq O(\mathrm{poly}(d) \cdot g/n).$$

(The exact dependence on maximum degree d is not mentioned in the paper.) Their proof relies on nontrivial results in the theory of Riemann surfaces, and it is not clear how to generalize it to the K_h-minor free case.

4.2 Upper bounding $\lambda_2'(G)$

It was Biswal, Lee, and Rao [BLR10] who successfully upper bounded $\lambda_2'(G)$ when G is K_h-minor free, resolving the conjecture by Spielman and Teng. Their approach deviates from the sphere embedding approach of previous works. Instead, they use results from l_1 embeddings of metric spaces, the flow/metric duality, and the relation between flow congestion and crossing number.

Theorem 4.3 (Upper bound on $\lambda_2'(G)$, [BLR10, Theorem 5.2, 5.3])**.** *Let $G = (V, E)$ be a graph with maximum degree d. Then,*

- *If G is of genus $g \geq 1$, then $\lambda_2'(G) \leq O(g \log^2 g \cdot \frac{d}{n})$.*[3]
- *If G is K_h minor free, then $\lambda_2'(G) \leq O(h^6 \log h \cdot \frac{d}{n})$.*

Step 1: Rayleigh quotient

First, we want to write the target quantity $\lambda_2'(G)$ as a Rayleigh quotient, then massage it so that all terms involved are distances, i.e. consisting only of terms $|f(u) - f(v)|$.

The following proposition is useful and admits a nice short proof.

Proposition 4.4. *If $f : V \to \mathbb{R}$ has mean zero, then $\sum_u |f(u)|^2 = \frac{1}{2n} \sum_{u,v} |f(u) - f(v)|^2$.*

Proof. Let X, Y be i.i.d. copies of f (with uniform distribution on V). Then $\mathbb{E}[X] = 0$. It is easy to check that $\sum_u |f(u)|^2 = n\mathbb{E}[X^2]$ and

$$\sum_{u,v} |f(u) - f(v)|^2 = n^2 \mathbb{E}[(X - Y)^2] = n^2(\mathbb{E}[X^2] + \mathbb{E}[Y^2] - 2\mathbb{E}[XY]) = 2n^2\mathbb{E}[X^2].$$

□

[3] In [BLR10] a weaker statement is proved, that $\lambda_2'(G) \leq O(g^3 \cdot \frac{d}{n})$. The bound in the theorem statement appears in [Kel+11] and relies upon a stronger result in [LS10] concerning the decomposition of shortest path metric on genus g graphs.

From the proposition, we have

$$\lambda_2'(G) = \min_{f \perp \mathbb{1}} \frac{\sum_{(u,v)\in E} |f(u)-f(v)|^2}{\sum_{u\in V} f(u)^2} = 2n \min_f \frac{\sum_{(u,v)\in E} |f(u)-f(v)|^2}{\sum_{u,v\in V} |f(u)-f(v)|^2},$$

where we drop the condition $f \perp \mathbb{1}$ in the last step because the quantity being minimized is translation invariant. We will focus on upper bounding $\min_f \frac{\sum_{(u,v)\in E} |f(u)-f(v)|^2}{\sum_{u,v\in V} |f(u)-f(v)|^2}$ from now on.

Step 2: From l_1 metric to shortest path metric

Next, we want to relax the minimization problem. We can write

$$\min_f \frac{\sum_{(u,v)\in E} |f(u)-f(v)|^2}{\sum_{u,v\in V} |f(u)-f(v)|^2} = \min_{\rho \in M_1(V)} \frac{\sum_{(u,v)\in E} \rho(u,v)^2}{\sum_{u,v\in V} \rho(u,v)^2} =: \min_{\rho\in M_1(V)} \mathcal{R}(\rho).$$

where the minimization is over all $\rho \in M_1(V)$, $M_1(V)$ denoting the set of all metrics on V induced by the l_1 distance on $\mathbb{R}$. The relaxation is by replacing $M_1(V)$ with $M_{sp}(V)$, the set of all vertex-weighted shortest path metrics on V. To see that the objective does not drop too much, it suffices to prove that, for any metric $\rho_s \in M_{sp}(V)$ there is a corresponding $\rho \in M_1(V)$, such that $\mathcal{R}(\rho)$ is not too much larger than $\mathcal{R}(\rho_s)$.

The theory of metric decomposition is useful here. We abstract out the theory and extract only the statement that we need. For details of the theory, refer to Section 4 of [BLR10].

For any $s : V \to \mathbb{R}_{>0}$, the vertex-weighted shortest path metric ρ_s induced by s is defined as

$$\rho_s(u,v) := \min\left\{ \sum_{w\in p} s(w) \,|\, p \text{ is a path from } u \text{ to } v \right\}.$$

Lemma 4.5 (Average distortion [BLR10, Theorem 4.4])**.** *For any graph $G = (V,E)$, there exists $\alpha(G) > 0$ such that the following holds: for any $p \geq 1$, and for any shortest path metric $\rho_s \in M_{sp}(V)$ there is a l_1 metric $\rho \in M_1(V)$, where $\rho(u,v) \leq \rho_s(u,v)$ for all $u,v \in V$ and*

$$\sum_{u,v\in V} \rho_s(u,v)^p \leq O_p(1)\cdot \alpha(G)^p \cdot \sum_{u,v\in V} \rho(u,v)^p.$$

Furthermore,

- *(Bartal [Bar96]) For any graph G, $\alpha(G) = O(\log n)$;*
- *(Lee, Sidiropoulos [LS10]) If G has genus $g \geq 1$, then $\alpha(G) = O(\log^2 g)$;*
- *(Klein, Plotkin, Rao [KPR93]) If G is K_h-minor free, then $\alpha(G) = O(h^2)$.*

Equipped with Lemma 4.5 and after a series of calculations, we sum up the work done in step 2 as follows:

Proposition 4.6 (λ_2' and metric spread)**.** *For any graph $G = (V,E)$, let $\alpha(G)$ be as defined in Lemma 4.5. Then,*

$$\lambda_2'(G) = 2n \min_{\rho\in M_0(V)} \mathcal{R}(\rho) \lesssim dn^3 \cdot \alpha(G)^2 \cdot \left[\max_{s:V\to\mathbb{R}_{>0}} \frac{\sum_{u,v\in V} \rho_s(u,v)}{\sqrt{\sum_{u\in V} s(u)^2}} \right]^{-2}.$$

The objective being maximized on RHS can be regarded as a concavification of

$$\frac{\sum_{u,v\in V}\rho_s(u,v)^2}{\sum_{u\in V}s(u)^2}$$

and will be useful in the next step. The proof is deferred to Appendix C.

Step 3: Metric spread and flow congestion

In this step, we relate the quantity

$$\max_{s:V\to\mathbb{R}_{>0}}\frac{\sum_{u,v\in V}\rho_s(u,v)}{\sqrt{\sum_{u\in V}s(u)^2}} =: \max_{s:V\to\mathbb{R}}\Lambda_s(G)$$

to a multicommodity flow problem (MFP) where the goal is to minimize some measure of congestion.

Given a flow solution F to a MFP, the congestion $c_F(u)$ at vertex $u\in V$ is the total amount of flow in F passing through u. The vertex p-congestion of F is defined as

$$con_p(F) := \left(\sum_{u\in V}c_F(u)^p\right)^{1/p}.$$

It turns out that metric spread maximization is strongly dual to flow 2-congestion minimization.

Lemma 4.7 (Flow/metric duality, [BLR10, Theorem 2.2])**.** *For any graph $G=(V,E)$,*

$$\min_F con_2(F) = \max_{s:V\to\mathbb{R}_{>0}}\Lambda_s(G),$$

where the minimum is taken over all flow solutions F on the uniform demand graph K_n (so that $D(u,v)=1$ for all $u\neq v\in V$).

The proof is by writing out the Lagrangian of the minimum congestion program, simplifying it to obtain $\Lambda_s(G)$, and lastly establishing a Slater point. Lemma 4.7 combined with Proposition 4.6 implies

$$\lambda_2'(G)\lesssim dn^3\cdot\alpha(G)^2\cdot\left(\min_F con_2(F)\right)^{-2}.$$

Step 4: Flow congestion and crossing number

It remains to *lower bound* the minimum 2-congestion $\min_F con_2(F)$, which will yield an upper bound on $\lambda_2'(G)$. The bounds obtained in [BLR10] are summarized below.

Lemma 4.8 (Congestion lower bound, [BLR10, Theorem 3.1, 3.11])**.** *Let $G=(V,E)$ be a graph and let F be a flow solution to the MFP with uniform demand graph K_n on G. Then,*

- *If G is of genus $g\geq 1$, and $n\gtrsim\sqrt{g}$, then $con_2(F)\gtrsim n^2/\sqrt{g}$;*
- *If G is K_h-minor free, and $n\gtrsim h\sqrt{\log h}+1$, then $con_2(G)\gtrsim n^2/(h\sqrt{\log h})$.*

The idea of proof for the genus g case is that, if $con_2(F)$ is too small, then we can round it to an integral flow F' with $con_2(F')$ as well, where F' being integral means each flow path carries an integral amount of flow. As the demand graph is K_n, this turns out to induce a drawing of K_n in a genus g surface $\mathcal{S}$ with few pairs of crossing edges (the graph G is embedded in $\mathcal{S}$, and edges of K_n are only allowed to cross at vertices of G), and contradict known lower bounds on the number of edge crossings when K_n is embedded in a genus g surface. The proof for the K_h-minor free case has the same spirit, but it proceeds by lower bounding the so-called intersection number instead of the crossing number. We will not delve into the details of the proof here, and remark that flow solutions can be rounded to integral flow solutions and relate to the combinatorial aspect of the graph.

Now we have all the ingredients for proving Theorem 4.3. Recall that at the end of step 3 we arrived at

$$\lambda_2'(G) \lesssim dn^3 \cdot \alpha(G)^2 \cdot \left(\min_F con_2(F)\right)^{-2}$$

where the minimum of F is taken over all flow solutions F on demand graph K_n. Applying Lemma 4.17 and Lemma 4.8, we have:

- When G is of genus g, $\alpha(G) \leq O(\log g)$ and $\min_F con_2(F) \geq n^2/\sqrt{g}$.
  Plugging these bounds in, we obtain
  $$\lambda_2'(G) \leq O(g \log^2 g \cdot d/n).$$
- When G is K_h-minor free, $\alpha(G) \leq O(h^2)$ and $\min_F con_2(F) \geq n^2/(h\sqrt{\log h})$.
  Plugging these bounds in, we obtain
  $$\lambda_2'(G) \leq O(h^6 \log h \cdot d/n).$$

This completes the proof of Theorem 4.3.

4.3 Reweighted eigenvalues

Before explaining how to obtain upper bounds on the reweighted eigenvalues $\lambda_2^*(G)$ and $\lambda_k^*(G)$, let us recall the definition of reweighted eigenvalues and state relevant results. Given a graph $G = (V, E)$ and a distribution $\pi : V \to \mathbb{R}$, the k-th reweighted eigenvalue can be regarded as an optimization problem: assign non-negative weights to the edges of G, so that the natural random walk P on the reweighted graph is π-reversible and the k-th eigenvalue of $I - P$ is maximized. We temporarily add self loops to the graph to ensure feasibility. Formally,

$$\begin{aligned}
\lambda_k^*(G,\pi) := \quad & \max_{P \geq 0} \quad && \lambda_k(I - P) \\
& \text{subject to} && \sum_v P(u,v) = 1 \quad \forall u \in V \\
& && P(u,v) = 0 \quad \forall (u,v) \notin E \\
& && \pi(u)P(u,v) = \pi(v)P(v,u) \quad \forall u,v \in V.
\end{aligned}$$

The $\lambda_2^*(G,\pi)$ program is convex – indeed it is a semidefinite program, and Roch [Roc05] showed that strong duality holds with an appropriately defined dual program. Define the k-dimensional dual program as

$$\begin{aligned}
\gamma^{(k)}(G,\pi) \quad & := \quad \min_{\substack{f:V\to\mathbb{R}^k \\ g:V\to\mathbb{R}_{\geq 0}}} \sum_{u\in V} \pi(u) g(u) \\
& \text{subject to} \quad g(u)+g(v) \geq \|f(u)-f(v)\|^2 \quad \forall (u,v)\in E \\
& \qquad \sum_{u\in V} \pi(u) f(u) = \vec{0} \\
& \qquad \sum_{u\in V} \pi(u) \|f(u)\|^2 = 1.
\end{aligned}$$

Theorem 4.9 (Strong duality, [Roc05])**.** *For any graph $G=(V,E)$ and any distribution $\pi : V \to \mathbb{R}$, $\lambda_2^*(G,\pi)$ and $\gamma^{(n)}(G,\pi)$ are strongly dual to one another.*

Define the π-weighted vertex expansion as

$$\psi(G,\pi) := \min_{S\subseteq V, 0<\pi(S)\leq 1/2} \frac{\pi(N(S))}{\pi(S)},$$

where for $T \subseteq V$, $\pi(T) := \sum_{u\in T} \pi(u)$. (Remark on notation: if we write $\lambda_k^*(G), \gamma^{(k)}(G), \psi(G)$ without the π parameter, we are considering the uniform distribution.) Recently, Olesker-Taylor and Zanetti [OZ22] proved a vertex Cheeger's inequality relating vertex expansion and reweighted second eigenvalue. The result has been subsequently improved by Kwok, Lau, and Tung [KLT22] and also Jain, Phong, and Vuong [JPV22], and we state the strongest form here.

Theorem 4.10 (Vertex Cheeger's inequality, [OZ22; KLT22; JPV22])**.** *For any graph $G=(V,E)$ with maximum degree d, and for any distribution $\pi : V \to \mathbb{R}$, if $\psi(G,\pi) \leq O(1)$ then*

$$\frac{\psi(G,\pi)^2}{\log d} \lesssim \lambda_2^*(G,\pi) \lesssim \psi(G,\pi).$$

The proof in [KLT22] consists of a projection step and a threshold rounding step. Refer to Section 3 of the paper for details.

Lemma 4.11 (Projection)**.** *For any graph $G=(V,E)$ with maximum degree d, and for any distribution $\pi : V \to \mathbb{R}$,*

$$\gamma^{(n)}(G,\pi) \leq \gamma^{(1)}(G,\pi) \lesssim \log d \cdot \gamma^{(n)}(G,\pi).$$

Lemma 4.12 (Threshold rounding)**.** *For any graph $G=(V,E)$ and any distribution $\pi : V \to \mathbb{R}$,*

$$\min(1, \psi(G,\pi))^2 \lesssim \gamma^{(1)}(G).$$

4.4 Upper bounding $\lambda_2^*(G,\pi)$

It turns out that straight-forward adaptations of proofs from [ST96] and [BLR10] suffice to obtain analogous upper bounds on $\lambda_2^*(G)$.

4.4.1 Recovering the planar separation theorem

The planar separation theorem states that every planar graph has a balanced vertex separator of size $O(\sqrt{n})$. We will apply a result from [ST96] on spherical cap embeddings of planar graphs, to recover the planar separation theorem in its full generality. A (closed) spherical cap $C \subseteq \mathbb{S}^2$ with center $z \in \mathbb{S}^2$ and radius $r > 0$ is the set of all points with geodesic distance at most r from the point z. Its area is $\Theta(r^2)$.

Lemma 4.13 (Spherical cap embedding of planar graphs, [ST96, Theorem 3.2, 4.2])**.** *For any planar graph $G = (V, E)$, there is a mapping $u \mapsto C(u)$ from $u \in V$ to spherical caps $C(u) \subset \mathbb{S}^2$ with center $z(u)$ and radius $r(u)$, such that:*

- *the interiors of the spherical caps do not overlap;*
- *$C(u)$ is tangent to $C(v)$ iff $(u, v) \in E$;*
- *$\sum_u z(u) = \vec{0}$.*

The observation is that the spherical cap embedding from Lemma 4.13 gives us a good solution to the $\gamma^{(3)}(G)$ program. Refer to Appendix C for the proof.

Lemma 4.14 (Dual objective upper bound)**.** *For any graph $G = (V, E)$, $\gamma^{(3)}(G) \leq O(1/n)$.*

Using the simple relaxation $\gamma^{(n)}(G) \leq \gamma^{(3)}(G)$ and applying Theorem 4.10 (vertex Cheeger's inequality), we would obtain $\psi(G) \leq O(\sqrt{(\log d)/n})$, where d is the maximum degree of the graph. However, we can do better. Project the $\gamma^{(3)}(G)$ solution to a $\gamma^{(1)}(G)$ solution by taking the best of the three coordinates, and we obtain $\gamma^{(1)}(G) \leq 3 \cdot \gamma^{(3)}(G)$. Now, applying Lemma 4.12 (threshold rounding), and noting that $\psi(G) \leq O(1)$, it follows that $\psi(G) \leq O(1/\sqrt{n})$.

Any subgraph $G[V']$, $V' \subseteq V$, of a planar graph $G = (V, E)$ is also planar, and the vertex expansion bound holds: $\psi(G[V']) \leq O(1/\sqrt{|V'|})$. By repeatedly removing vertex cuts $S' \subseteq V'$ with $\psi(S') \leq O(1/\sqrt{|V'|})$, we obtain a balanced vertex cut of size $O(\sqrt{n})$, and the planar separation theorem follows.

4.4.2 Adaptation of [BLR10] to upper bound $\lambda_2^*(G)$

The proof in [BLR10] can be modified to upper bound $\lambda_2^*(G)$ for special classes of graphs G. Starting with the proof presented in Section 4.2, the main changes needed are to work with a different Rayleigh quotient and to use π-weighted versions of various intermediate results.

Theorem 4.15 (Upper bound on $\lambda_2^*(G, \pi)$)**.** *Let $G = (V, E)$ be a graph and π be any distribution on V. Let $\pi_{\max} := \max_{u \in V} \pi(u)$. Then,*

- *If G is of genus $g \geq 1$, then $\lambda_2^*(G, \pi) \leq O(\pi_{\max} \cdot g \log^2 g)$.*
- *If G is K_h-minor free, then $\lambda_2^*(G, \pi) \leq O(\pi_{\max} \cdot h^6 \log h)$.*

In particular, if π is the uniform distribution, then we obtain $\lambda_2^*(G) \leq O(g \log^2 g/n)$ in the genus g case and $\lambda_2^*(G) \leq O(h^6 \log h/n)$ in the K_h-minor free case. They are similar in form to the bounds on $\lambda_2'(G)$, except that there is no dependence on maximum degree d. The proof is as follows:

I. Modifications to steps 1-3 of Section 4.2

In step 1, we need to generalize Proposition 4.4 to arbitrary distribution π:

Proposition 4.16 (Generalization of Proposition 4.4)**.** *Let $\pi : V \to \mathbb{R}$ be a distribution on V. If $f : V \to \mathbb{R}$ satisfies $\sum_{u \in V} \pi(u) f(u) = 0$, then*

$$\sum_u \pi(u)|f(u)|^2 = \frac{1}{2} \sum_{u,v} \pi(u)\pi(v)|f(u) - f(v)|^2.$$

The proof is identical to that of Proposition 4.4.

In step 2, the result about average distortion of l_1 line embeddings in fact works for arbitrary distribution π. The interested reader can look up [FHL08, Appendix A.2] and [LS10, Theorem 4.2]. (The keyword is "product weight", weights of the form $\omega(u, v) = \pi(u)\pi(v)$.)

Lemma 4.17 (Generalization of Lemma 4.5, [FHL08; LS10])**.** *Under the same assumptions as Lemma 4.5, for any shortest path metric $\rho_s \in M_{sp}(V)$ there is a l_1 metric $\rho \in M_1(V)$, where $\rho(u, v) \le \rho_s(u, v)$ for all $u, v \in V$ and*

$$\sum_{u,v \in V} \pi(u)\pi(v)\rho_s(u, v)^p \le O_p(1) \cdot \alpha(G)^p \cdot \sum_{u,v \in V} \pi(u)\pi(v)\rho(u, v)^p.$$

Furthermore, we have the same upper bounds on $\alpha(G)$ as in Lemma 4.5.

Combining Proposition 4.16 and Lemma 4.17, and after some calculations which are deferred to Appendix C, we arrive at this checkpoint:

Proposition 4.18 (λ_2^* and metric spread)**.** *For any graph $G = (V, E)$ and distribution π on V, let $\alpha(G)$ be as defined in Lemma 4.17. Then, for any π-reversible reweighting P of G (i.e. P satisfies the constraints of the $\lambda_2^*(G, \pi)$ program as defined in Section 4.3),*

$$\lambda_2(I - P) = 2 \min_{\rho \in M_0(V)} \frac{\sum_{(u,v) \in E} \pi(u)P(u, v)\rho(u, v)^2}{\sum_{u \in V} \pi(u)\pi(v)\rho(u, v)^2} \lesssim \alpha(G)^2 \cdot \left[\max_{s:V \to \mathbb{R}_{>0}} \frac{\sum_{u,v \in V} \pi(u)\pi(v)\rho_s(u, v)}{\sqrt{\sum_{u \in V} \pi(u)s(u)^2}} \right]^{-2}.$$

As a consequence,

$$\lambda_2^*(G, \pi) \lesssim \alpha(G)^2 \cdot \left[\max_{s:V \to \mathbb{R}_{>0}} \frac{\sum_{u,v \in V} \pi(u)\pi(v)\rho_s(u, v)}{\sqrt{\sum_{u \in V} \pi(u)s(u)^2}} \right]^{-2}.$$

Compared with Proposition 4.6 (λ_2' and metric spread), the most notable feature is that this upper bound has no dependence on d. This is reflected in the lack of d-dependence in Theorem 4.15 (upper bound on $\lambda_2^*(G)$). The apparent lack of polynomial dependence on n on the RHS is due to the π coefficients. Indeed, if we substitute $\pi(u) = 1/n$ and pull out the n terms from the fraction, the n^3 dependence resurfaces.

The next modification needed is a π-weighted flow/metric duality in step 3. Define π-weighted vertex p-congestion of a flow solution F as

$$con_p(F, \pi) := \left(\sum_{u \in V} \pi(u)^{-1} c_F(u)^p \right)^{1/p},$$

and define[4]

$$\Lambda_s(G,\pi) := \frac{\sum_{u,v\in V}\pi(u)\pi(v)\rho_s(u,v)}{\sqrt{\sum_{u\in V}\pi(u)s(u)^2}}.$$

This time, the corresponding flow problem is no longer uniform. The demand between u and v should be $\pi(u)\pi(v)$ instead of 1. We use $(K_n, \pi\times\pi)$ to denote the π-product demand graph, where the demand between u and $v \neq u$ is $D(u,v) = \pi(u)\pi(v)$. Roughly speaking, if $\pi(u)$ is small, we need to send fewer units of flow from u to other vertices, and so we may be inclined to lower the capacity of u, and weight the congestion (which is amount of flow per unit of capacity) at u by $\pi(u)^{-1}$. This partly explains the slightly odd-looking form that $con_p(F,\pi)$ takes, although the most accurate explanation would be by taking the dual of $\max_s \Lambda_s(G,\pi)$.

The proof of the following result is almost identical to that of Lemma 4.7 and will be omitted.

Lemma 4.19 (Generalization of Lemma 4.7)**.** *For any graph $G=(V,E)$ and distribution π on V,*

$$\min_F con_2(F,\pi) = \max_{s:V\to\mathbb{R}_{>0}} \Lambda_s(G,\pi),$$

where the minimum is taken over all flow solutions on the π-product demand graph $(K_n,\pi\times\pi)$.

Let $\pi_{\max} := \max_{u\in V}\pi(u)$. Lower bounding each $\pi(u)^{-1}$ by $\pi_{\max}^{-1}$ in the expression $con_2(F,\pi)$, we arrive at

$$\lambda_2^*(G) \lesssim \alpha(G)^2 \cdot \left(\min_F con_2(F,\pi)\right)^{-2} \lesssim \pi_{\max}\cdot\alpha(G)^2\cdot\left(\min_F con_2(F)\right)^{-2},$$

where again $\min_F$ is over all flow solutions on the demand graph $(K_n,\pi\times\pi)$. We remind the reader that

$$con_2(F) = \left(\sum_{u\in V} c_F(u)^2\right)^{1/2}$$

is an expression with no π-dependence, although F depends on π.

II. Congestion lower bound for product demand

Finally, in step 4, we need a lower bound on the 2-congestion $con_2(F)$ for π-product demand graphs. This part is the most interesting. First of all, note that if $\pi_{\max}\geq 1/2$ then the result of Theorem 4.15 trivially holds; so, assume $\pi_{\max} < 1/2$. We also assume WLOG that $\pi(u)>0$ for all $u\in V$, because if $\pi(u)=0$ we can drop u from the graph without affecting anything.

The idea is to round our given weight function π to a *rational* weight function π_q (we emphasize that the entries of a weight function may no longer sum to 1), then to an *integral* weight function π_z. Then, we consider an auxiliary graph $G'=(V',E')$, such that the congestion of solution F' to the MFP with unit demand graph $K_{n'}$ on G' is closely related to the congestion of solution F to the MFP with π_z-product demand graph $(K_n,\pi_z\times\pi_z)$ on G. This enables us to apply the results of [BLR10] and wrap up the proof.

A weight function $\pi: V\to\mathbb{R}_{>0}$ is said to be *rational* if $\pi(u)\in\mathbb{Q}_{>0}$ for all $u\in V$. In other words, there exists $M\in\mathbb{Z}_{>0}$ such that $M\pi(u)$ is a positive integer for all $u\in V$. π is said to be *integral* if $\pi(u)\in\mathbb{Z}_{>0}$ for all $u\in V$.

[4]Note that when π is the uniform distribution, $\Lambda_s(G,\pi)$ is off by a factor of $n^{3/2}$ from $\Lambda_s(G)$ defined in Section 4.2.

Proposition 4.20 (Rounding to rational weight function)**.** *Given any weight function $\pi : V \to \mathbb{R}_{>0}$, there exists a weight function $\pi_q : V \to \mathbb{Q}_{>0}$ such that $\pi(u) \leq \pi_q(u) \leq 2\pi(u)$ for all $u \in V$.*

The proof is deferred to Appendix C.

The implication is that we can replace the π-product demand graph $(K_n, \pi \times \pi)$ by a rational π_q-product demand graph $(K_n, \pi_q \times \pi_q)$. For a flow solution F to $(K_n, \pi \times \pi)$, $4F$ is a flow solution to $(K_n, \pi_q \times \pi_q)$, and so it suffices to lower bound the 2-congestion $con_2(F_q)$ over flow solutions F_q to $(K_n, \pi_q \times \pi_q)$.

Now, fix $M \in \mathbb{Z}_{>0}$ such that $M\pi_q(u)$ is a positive integer for all $u \in V$. Let $\pi_z := M\pi_q$. Then π_z is an integral weight function. It is clear that there is a 1-1 correspondence between flow solutions to $(K_n, \pi_q \times \pi_q)$ and solutions to $(K_n, \pi_z \times \pi_z)$, via $F_q \leftrightarrow F_z := M^2 \cdot F_q$. Moreover, $con_2(F_z) = M \cdot con_2(F_q)$. Hence, it suffices to lower bound the 2-congestion $con_2(F_z)$ over flow solutions F_z to $(K_n, \pi_z \times \pi_z)$, where π_z is an integral weight function.

We now describe the key construction.

Definition 4.21 (Spike graph)**.** *Given a graph $G = (V, E)$ and an integral weight function $\pi_z : V \to \mathbb{Z}_{>0}$, the spike graph*[5] *$G' = G'(G, \pi_z) = (V', E')$ of G with respect to π_z is constructed as follows:*

- *For each $u \in V$, there are $\pi_z(u)$ copies of u in V': $u_0, u_1, \ldots, u_{\pi_z(u)-1}$. u_0 can be considered the original/central vertex;*
- *For each $(u, v) \in E$, there is an edge $(u_0, v_0) \in E'$;*
- *For each $u \in V$ and $1 \leq i < \pi_z(u)$, there is an edge $(u_0, u_i) \in E'$.*

Refer to Figure 2 for an example.

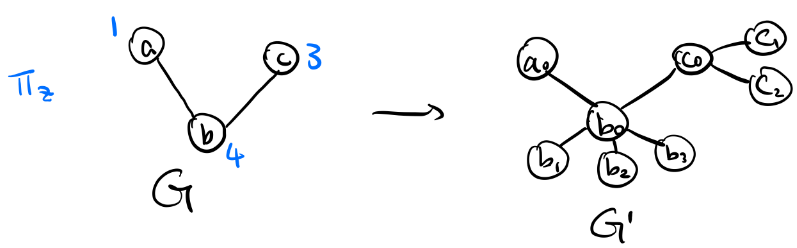

Figure 2: On the left is a graph $G = (V, E)$ equipped with an integral weight function π_z on V. On the right is the associated spike graph G'.

We now relate the flow problem $K_{n'}$ on the spike graph G' with the flow problem $(K_n, \pi_z \times \pi_z)$ on G. The proof is deferred to Appendix C.

Lemma 4.22 (Spike graph congestion)**.** *There exists universal constants $A, B > 0$ such that*

$$A \cdot \min_{F'} con_2(F') \leq \min_{F_z} con_2(F_z) \leq B \cdot \min_{F'} con_2(F'),$$

where the minimum of F_z is taken over all flow solutions F_z with demand graph $(K_z, \pi_z \times \pi_z)$ and the minimum of F' is taken over all flow solutions F' with demand graph $K_{n'}$ on the spike graph G' of G with respect to π_z.

[5] We are not sure if this construction has already been named in the literature. We call it spike graph because it looks like at each vertex there are spikes growing out of it.

Before we apply the flow congestion lower bound in Lemma 4.8, we must check that the topological properties (genus g, K_h-minor free) of G' are the same as that of G. Again, the proof can be found in Appendix C.

Lemma 4.23 (Spike graph topology)**.** *Given a graph $G = (V, E)$ and an integral weight function $\pi_z : V \to \mathbb{Z}_{>0}$, let G' be the spike graph $G' = (V', E')$ of G with respect to π_z. Then,*

- *If G has genus g, then G' has genus g;*
- *If G is K_h-minor free (for $h \geq 3$), then G' is K_h-minor free.*

Now we are ready to prove Theorem 4.15. Recall that we have shown

$$\lambda_2^*(G) \lesssim \pi_{\max} \cdot \alpha(G)^2 \cdot \left(\min_F con_2(F)\right)^{-2},$$

where the minimum of F is taken over all flow solutions F on demand graph $(K_n, \pi \times \pi)$. By the $\pi \mapsto \pi_q \mapsto \pi_z$ transformation, we know that

$$\min_F con_2(F) \geq \frac{1}{4M^2} \min_{F_z} con_2(F_z),$$

where the minimum of F_z is taken over all flow solutions F_z on demand graph $(K_n, \pi_z \times \pi_z)$.

The number of vertices n' of the spike graph G' is at least

$$n' = \sum_{u \in V} \pi_z(u) = \sum_{u \in V} \lceil M\pi(u) \rceil \geq M.$$

Applying Lemma 4.17 and Lemma 4.8, (note that $n' \geq M \geq n/2$, so whenever n satisfies the conditions $n \gtrsim \sqrt{g}$ or $n \gtrsim h\sqrt{\log h} + 1$, so does n'), we have:

- When G is of genus g, $\alpha(G) \leq O(\log g)$ and

  $$\min_F con_2(F) \geq \frac{1}{4M^2} \min_{F_z} con_2(F_z) \gtrsim \frac{1}{M^2} \min_{F'} con_2(F') \gtrsim \frac{1}{M^2} \cdot \frac{(n')^2}{\sqrt{g}} \geq 1/\sqrt{g}.$$

  Plugging these bounds in, we obtain

  $$\lambda_2^*(G, \pi) \leq O(\pi_{\max} \cdot g \log^2 g).$$

- When G is K_h-minor free, $\alpha(G) \leq O(h^2)$ and

  $$\min_F con_2(F) \geq \frac{1}{4M^2} \min_{F_z} con_2(F_z) \gtrsim \frac{1}{M^2} \cdot \frac{(n')^2}{h \log h} \geq 1/(h\sqrt{\log h}).$$

  Plugging these bounds in, we obtain

  $$\lambda_2^*(G, \pi) \leq O(\pi_{\max} \cdot h^6 \log h).$$

This completes the proof of Theorem 4.15.

Immediately by applying Theorem 4.10 (vertex Cheeger's inequality), we obtain the following bounds on $\psi(G, \pi)$ for the same classes of graph G:

Corollary 4.24 (Upper bounds on $\psi(G,\pi)$). *Let $G=(V,E)$ be a graph of maximum degree d. Let π be a distribution on V and let $\pi_{max} := \max_{u\in V}\pi(u)$. Then,*

- *If G is of genus $g \geq 1$, then $\psi(G,\pi) \leq O(\sqrt{\log d \cdot \pi_{\max} \cdot g \log^2 g})$.*
- *If G is K_h-minor free, then $\psi(G,\pi) \leq O(\sqrt{\log d \cdot \pi_{\max} \cdot h^6 \log h})$.*

Note that when π is uniform, these bounds are not the best obtained so far. In fact, in [BLR10] they used their approach to provide direct bounds on $\psi(G)$. The bounds they obtained are better than the ones in Corollary 4.24 by a factor of $\sqrt{\log d}$. Refer to Theorem 5.5 of their paper for more details.

5 Summary

There are several results presented in this report that are new to the best of our knowledge:

- Theorem 3.8 (cut matching game for $\psi(G)$) which gives an $O(\log n)$-approximation of $\psi(G)$ with runtime $O(m^{1+o(1)})$;
- Lemma 4.14 which upper bounds $\gamma^{(3)}(G)$ by $O(1/n)$ and hence $\psi(G)$ by $O(1/\sqrt{n})$ for planar graphs G, providing an alternative proof of the planar separation theorem;
- Theorem 4.15 which upper bounds $\lambda_2^*(G,\pi)$ for G of genus $g \geq 1$ and for G being K_h-minor free. The lower bound on $con_2(F)$ for π-product demand graphs generalizes 4.8 and may be of independent interest.

There are some unfortunate omissions in this project pertaining to the theme of using flows in graph expansion parameters. They are no less important than the results presented in the report, but they had to be excluded due to time and space constraint as well as a change of focus on presenting the new results. Nevertheless, we feel compelled to do them justice by mentioning them here, near the end of the report:

- $O(\log n)$-approximation of $\phi'(G)$ by Leighton and Rao [LR99]. Their idea is to prove an approximate max flow/min cut theorem for uniform MFP's: $\mathcal{C}/\log n \lesssim \mathcal{F} \leq \mathcal{C}$. Since min cut for the uniform MFP is exactly the sparsest cut, which is essentially the same as edge expansion $\phi'(G)$, computing max flow for uniform MFP can yield an $O(\log n)$ approximation of $\phi'(G)$.
- $O(\sqrt{\log n})$-approximation of $\phi'(G)$ by Arora, Rao, and Vezirani [ARV09]. The second approach in their paper is called "expander flows" and the idea is to consider MFP's where the demand graph could be an expander, instead of K_n as in [LR99].
- Upper bounds on $\lambda_k'(G)$ for G planar, genus g, K_h-minor free. These are the results by Kelner et al. [Kel+11]. While similar in large to the approach in [BLR10], they introduced a new concept called "subset flows" to deal with k-partitions of metric spaces. It looks promising to extend their approach to give upper bounds on $\lambda_k^*(G)$, or even $\lambda_k^*(G,\pi)$, for these same classes of graphs. In fact, we already know that their approach to upper bound $\lambda_k'(G)$ (i.e. uniform distribution) can be easily modified to obtain an upper bound on $\lambda_k^*(G)$, indeed a bound of the same form except with the maximum degree dependence removed.

To sum up, we have seen how flow techniques were used in approximating and upper bounding graph expansion parameters. There are a few recurring themes:

- Max flow/min cut duality. We have seen how the classical (single commodity) max flow/min cut duality has been used time and time again to construct, from a flow solution, a cut with matching value – for example, in the proofs of Lemma 2.3 and Lemma 3.2. The results by [LR99] and [ARV09] can also be interpreted as an approximate max flow/min cut result for multicommodity flows.

- Expansion certification using demand graph of MFP. We have seen that, if a demand graph H of a MFP has good expansion and can be embedded in a graph G with low congestion, then G must have good expansion as well. Many approximation algorithms, for example the cut-matching game, are based on this theme.

- Efficiency of flow computations. It has been said that flow computations are the "golden standard" of approximation algorithms. If an approximation algorithm has runtime dominated by few s-t flow computations, then it is truly fast. The cut-matching game is devised with this principle in mind. Some other follow-up work of [ARV09], for example [AHK10] and [She09], improved the runtime of [ARV09] using approximate multicommodity flow computations and s-t flow computations, both of which are very efficient.

- Flow/metric duality. We have seen the theme of flow/metric duality in Section 4, specifically that the maximization of $\Lambda_s(G, \pi)$ – the spread of vertex weighted shortest-path metrics – is dual to the minimization of the 2-congestion of π-product MFP's. This duality is the bridge between the spectrum and the combinatorics of the graph. Although not mentioned here, already in [LR99] is this theme utilized in approximating $\phi'(G)$. They take the dual of the max flow (of the uniform MFP) and obtain a shortest-path metric program where the goal is to maximize the total distance between the commodities. The expander flow approach in [ARV09] also has a similar dual interpretation, and indeed the dual corresponds to the first approach in their paper of SDP rounding.

Once again, for Theorem 3.8 (Cut matching agme for $\psi(G)$), I thank Robert Wang for working out the details together and Yueheng Zhang and Lap Chi Lau for the discussions. I thank also Thatchaphol Saranurak for confirming the correctness of the result.

Thank you for reading the report and hope you enjoyed it! :)

A Deferred proofs of Section 2

Proof of Lemma 2.2 (Truncation). Let's prove $R'(f_+) \leq \lambda_2'$; the result about $R'(f_-)$ is analogous.

By definition, f satisfies $L'f = \lambda_2' f$, and when written out explicitly we get

$$\sum_{v:(u,v)\in E} (f(u) - f(v)) = \lambda_2' f(u) \quad \forall u \in V.$$

Let $V^+ := \{u \in V : f(u) > 0\}$. Multiplying each equation by $f(u)$ and summing over $u \in V^+$, we get

$$\sum_{u\in V^+} \sum_{v:(u,v)\in E} f(u)(f(u) - f(v)) = \lambda_2' \sum_{u\in V^+} f(u)^2.$$

For each edge (u, v) appearing in LHS, if $u, v \in V^+$ then there is a $f(u)(f(u)-f(v))$ term and a $f(v)(f(v)-f(u))$ term, and they sum to $(f(u) - f(v))^2$. Otherwise, if $u \in V^+$ and $v \notin V^+$, then there is only a $f(u)(f(u) - f(v))$ term which is greater than $f(u)^2 = (f(u) - 0)^2$. Therefore,

$$\sum_{(u,v)\in E} (f_+(u) - f_+(v))^2 \leq \sum_{u\in V^+} \sum_{v:(u,v)\in E} f(u)(f(u) - f(v)) = \lambda_2' \sum_{u\in V^+} f(u)^2 = \lambda_2' \sum_{u\in V} f_+(u)^2,$$

and we are done after rearranging. □

Proof of Lemma 2.3 (Full flow). We want to prove that the flow problem has value $(1+\psi)|A|$. Suppose not. Then by max flow/min cut theorem, we can find an s-t edge cut with total cost less than $(1 + \psi)|A|$. In particular, some edges from s to $u \in A$ are not cut.

Let $S \subseteq A$ be the set of vertices in A that remain directly connected from u. Note that $0 < |S| \leq |A| \leq n/2$. We want to show that $\psi(S) < \psi$ which contradicts the definition of $\psi = \psi(G)$. In order to disconnect s from t, we need to disconnect S from t. Check that the vertices in B (on the t side) directly connected to S are precisely those vertices in $S \cup N(S) \subseteq B$. For each $v \in S \cup N(S)$, we need to delete at least one edge incident to it, each of which has cost 1. Remember also that we have cut the edges from s to u for $u \in A \setminus S$, and each has cost $(1 + \psi)$. Therefore, the value of the cut is at least

$$(1 + \psi)(|A| - |S|) + 1 \cdot |S \cup N(S)| = (1 + \psi)(|A| - |S|) + (|S| + \psi(S) \cdot |S|) = (1 + \psi)|A| + (\psi(S) - \psi)|S|,$$

and for it to be strictly less than $(1+\psi)|A|$, we must have $\psi(S) < \psi$, and the desired contradiction follows. □

Proof of Lemma 2.4 (Flow inequality I). It suffices to show that

$$\sum_{(u,v)\in E} h(u, v)^2 (g(u)^2 + g(v)^2) \leq (2 + \psi^2) \sum_{u\in V} g(u)^2.$$

The coefficient of $g(u)^2$ on the LHS is

$$\sum_{v:(u,v)\in E} h(u, v)^2 + \sum_{v:(v,u)\in E} h(v, u)^2.$$

We just need to show that this is at most $(2+\psi^2)$. Since $0 \le h \le 1$, $\sum_{v:(u,v)\in E} h(u,v) \le 1+\psi$, and $\sum_{v:(v,u)\in E} h(v,u) \le 1$, we conclude that

$$\sum_{v:(u,v)\in E} h(u,v)^2 + \sum_{v:(v,u)\in E} h(v,u)^2 \le (1+\psi^2) + 1 = 2+\psi^2,$$

and the proof is complete. $\square$

Proof of Lemma 2.5 (Flow inequality II). We can write

$$\sum_{(u,v)\in E} h(u,v)(g(u)^2 - g(v)^2) = \sum_{u\in A}\sum_{v\in B} h(u,v)(g(u)^2 - g(v)^2),$$

by adding the terms $h(u,u)(g(u)^2 - g(u)^2)$ to LHS. Then,

$$\sum_{u\in A}\sum_{v\in B} h(u,v)(g(u)^2 - g(v)^2) = (1+\psi)\sum_{u\in A} g(u)^2 - \sum_{v\in B}\left(\sum_{u\in A} h(u,v)\right) g(v)^2 \ge \psi \sum_{u\in V} g(u)^2.$$

The last inequality holds because $\sum_{u\in A} h(u,v) \le 1$ and $g(v)^2 = 0$ for $v \notin A$. $\square$

B Deferred proofs of Section 3

Proof of Lemma 3.2 (Edge cut and congestion). The proof largely follows that presented in [KRV09]. For convenience let's drop the subscripts i.

Let $\varepsilon > 0$. Consider the s-t flow problem with vertex capacity $(c-\varepsilon)$. Then, we can only send $f < n/2$ units of flow from s to t. Let's find a cut (of edges) with cut value f.

Define the following vertex subsets:

- Let $X^e \subseteq X$ be the set of vertices u on the X side such that (s,u) gets cut; similar for Y^e;
- Let $X^r \subseteq X$ be the remaining vertices in X, i.e. $X^r = X - X^e$; similar for Y^r;
- Let X^+ be the set of vertices connected to s after the cut; similar for Y^+.

Figure 3 shows a sample flow network with the corresponding subsets X^e, Y^e, X^+, Y^+.

(We crucially need $f < n/2$ to make sure neither X^+ nor Y^+ is empty, hence the ε trick.)

One of the sets X^+ and Y^+ will be chosen as S_i in the end. Note that:

- $X^r \subseteq X^+$ and $Y^r \subseteq Y^+$;
- X^+ and Y^+ are disjoint;
- We can WLOG assume $X^+ \cup Y^+ = V$. This is because if there is a vertex u that is disconnected from both s and t, then either the graph is disconnected (in which case $\phi'(G) = 0$), or we can add back some edge incident to u. The cut value decreases, and t remains disconnected from s.

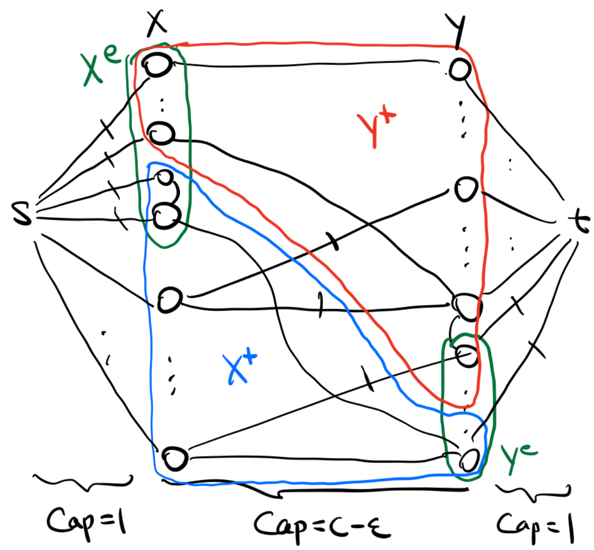

Figure 3: A flow network with the corresponding subsets X^e, Y^e, X^+, Y^+.

Because of the final observation, all the cut edges are either between u and X^e, between v and Y^e, or between X^+ and Y^+.

Now let's bound $\phi'(X^+)$. Cut value being equal to f means

$$(c-\varepsilon) \cdot |E(X^+, Y^+)| + (|X^e| + |Y^e|) = f \leq n/2.$$

Since $|X^+| \geq |X^r| = n/2 - |X^e|$, we have

Therefore,

$$\begin{aligned} \phi'(X^+) = \frac{|E(X^+,Y^+)|}{|X^+|} &\leq \frac{(c-\varepsilon)^{-1} \cdot (n/2 - (|X^e| + |Y^e|))}{n/2 - |X^e|} \\ &\leq \frac{(c-\varepsilon)^{-1} \cdot (n/2 - |X^e| - |Y^e|)}{n/2 - |X^e| - |Y^e|} \\ &= 1/(c-\varepsilon). \end{aligned}$$

Similarly we get $\phi'(Y^+) \leq 1/(c-\varepsilon)$. Since X^+ and Y^+ are disjoint, one of them will have size $\leq n/2$, and letting S_i to be that set, we have $|S_i| \leq n/2$ and $\phi'(S_i) \leq 1/(c_i - \varepsilon)$. Finish by letting $\varepsilon \to 0$. □

Proof of Lemma 3.5 (Vertex cut and congestion). The proof is quite similar to that of Lemma 3.2. For convenience let's drop the subscripts i.

Let $\varepsilon > 0$. Consider the s-t flow problem with vertex capacity $(c-\varepsilon)$. Then, we can only send $f < n/2$ units of flow from s to t. Let's find a cut (of edges $(s,u), (v,t)$ and of vertices) with cut value f.

Define the following vertex subsets:

- let $X^v \subseteq X$ be the set of vertices on the X side that gets cut; similar for Y^v;

- let $X^e \subseteq X$ be the set of vertices u on the X side such that (s,u) gets cut; similar for Y^e;
- let $X^r \subseteq X$ be the remaining vertices in X, i.e. $X^r = X - X^v - X^e$; similar for Y^r;
- let X^+ be the set of vertices connected to s after the cut; similar for Y^+.

(We crucially need $f < n/2$ to make sure neither X^+ nor Y^+ is empty, hence the ε trick.)

One of the sets X^+ and Y^+ will be chosen as S_i in the end. Note that:

- $X^r \subseteq X^+$ and $Y^r \subseteq Y^+$;
- $N(X^+) \subseteq X^v \cup Y^v$ and $N(Y^+) \subseteq X^v \cup Y^v$. This is because, once X^v and Y^v are removed, there is no more path from X^r to Y^r, and as X^+ is reachable from X^r and Y^+ is reachable from Y^r, this means there is no more path from X^+ to Y^+;
- X^+ and Y^+ are disjoint.

Now let's bound $\psi(X^+)$. Cut value being equal to f means

$$(c-\varepsilon)(|X^v|+|Y^v|)+(|X^e|+|Y^e|) = f \le n/2.$$

Rearranging and from $N(X^+) \subseteq X^v \cup Y^v$ we get

$$|N(X^+)| \le |X^v|+|Y^v| \le \frac{1}{(c-\varepsilon)}(n/2-(|X^e|+|Y^e|)).$$

Now let's lower bound $|X^+|$. We have:

$$|X^+| \ge |X^r| = n/2 - |X^v| - |X^e| \ge n/2 - (|X^v|+|Y^v|) - (|X^e|+|Y^e|).$$

It will actually be a bit more convenient to lower bound $\psi(X^+)^{-1}$:

$$\begin{aligned}
\psi(X^+)^{-1} = \frac{|X^+|}{|N(X^+)|} &\ge (c-\varepsilon)\cdot\frac{n/2-(|X^v|+|Y^v|)-(|X^e|+|Y^e|)}{n/2-(|X^e|+|Y^e|)} \\
&\ge (c-\varepsilon)\cdot\left(1-\frac{|X^v|+|Y^v|}{n/2-(|X^e|+|Y^e|)}\right) \\
&\ge c-\varepsilon-1 \\
&\quad \left(\because |X^v|+|Y^v| \le \frac{1}{(c-\varepsilon)}(n/2-(|X^e|+|Y^e|))\right).
\end{aligned}$$

Similarly we get $\psi(Y^+)^{-1} \ge c-\varepsilon-1$. Since X^+ and Y^+ are disjoint, one of them will have size $\le n/2$, and letting S_i to be that set, we have $|S_i| \le n/2$ and $\psi(S_i) \le 1/(c_i-1-\varepsilon)$. Finish by letting $\varepsilon \to 0$. □

Proof of Lemma 3.7 (Certifying vertex expansion). Since the MFP with demand graph M_i can be embedded in G with vertex congestion c_i, adding up the flow solutions, the MFP with demand graph $H_T := \sqcup_{i\in T} M_i$ can be embedded in G with vertex congestion at most $\sum_i c_i$.

Now, let $S \subseteq V$ such that $|S| \le n/2$. We want to show that

$$\psi(S) \ge \phi'(H_T)/\sum_i c_i.$$

From set S, H_T demands that we send at least $\phi'(H_T)\cdot|S|$ units of flow out of S. Each outgoing flow path must go through $N(S)$, so the total congestion over all vertices in $N(S)$ is at least $\phi'(H_T)\cdot|S|$. The feasibility of the flow problem when vertex capacity is $\sum_i c_i$ implies that

$$|N(S)|\cdot\sum_i c_i \geq \phi'(H_T)\cdot|S|.$$

Rearranging and minimizing over S, we are done. $\square$

C Deferred proofs of Section 4

Proof of Proposition 4.6 (λ_2' and metric spread). The equality directly follows from 4.4. The inequality follows from a series of calculations:

$$\begin{aligned}
\min_{\rho\in M_1(V)}\mathcal{R}(\rho) &= \min_{\rho\in M_1(V)}\frac{\sum_{(u,v)\in E}\rho(u,v)^2}{\sum_{u,v\in V}\rho(u,v)^2}\\
&\lesssim \min_{s:V\to\mathbb{R}_{>0}}\alpha(G)^2\frac{\sum_{(u,v)\in E}\rho_s(u,v)^2}{\sum_{u,v\in V}\rho_s(u,v)^2} &&\text{(Lemma 4.5)}\\
&= \alpha(G)^2\cdot\min_{s:V\to\mathbb{R}_{>0}}\frac{\sum_{(u,v)\in E}(s(u)+s(v))^2}{\sum_{u,v\in V}\rho_s(u,v)^2}\\
&\leq \alpha(G)^2\cdot\min_{s:V\to\mathbb{R}_{>0}}\frac{\sum_{u\in V}2d\cdot s(u)^2}{\sum_{u,v\in V}\rho_s(u,v)^2}\\
&\leq 2dn^2\cdot\alpha(G)^2\cdot\min_{s:V\to\mathbb{R}_{>0}}\frac{\sum_{u\in V}s(u)^2}{(\sum_{u,v\in V}\rho_s(u,v))^2} &&\text{(Cauchy-Schwarz)}\\
&\lesssim dn^2\cdot\alpha(G)^2\cdot\left[\max_{s:V\to\mathbb{R}_{>0}}\frac{\sum_{u,v\in V}\rho_s(u,v)}{\sqrt{\sum_{u\in V}s(u)^2}}\right]^{-2}.
\end{aligned}$$

$\square$

Proof of Lemma 4.14 (Dual objective upper bound). Take the spherical cap embedding $u\mapsto C(u)$, the cap $C(u)$ having center $z(u)$ and radius $r(u)$, from Lemma 4.13. Let $A(u)$ be the area of $C(u)$; recall that $A(u)=\Theta(r(u)^2)$. We construct a feasible solution to the $\gamma^{(3)}(G)$ program based on the embedding.

Take $f(u):=z(u)$ and $g(u):=2r(u)^2$. Check that $\sum_{u\in V}\frac{1}{n}f(u)=\vec{0}$, $\sum_{u\in V}\frac{1}{n}\|f(u)\|^2=1$, and

$$g(u)+g(v)=2(r(u)^2+r(v)^2)\geq(r(u)+r(v))^2\geq\|f(u)-f(v)\|^2\quad\forall(u,v)\in E.$$

Therefore, (f,g) is a feasible solution, and its objective is

$$\sum_{u\in V}\frac{1}{n}g(u)=\frac{2}{n}\sum_{u\in V}r(u)^2\lesssim\frac{1}{n}\sum_{u\in V}A(u)\leq\frac{1}{n}(\text{Area of }\mathbb{S}^2)\leq O(1/n).$$

$\square$

Proof of Proposition 4.18 (λ_2^* and metric spread). The equality directly follows from Proposition 4.16. The inequality follows from a series of calculations:

$$\begin{aligned}
& \min_{\rho \in M_1(V)} \frac{\sum_{(u,v)\in E} \pi(u)P(u,v)\rho(u,v)^2}{\sum_{u,v\in V} \rho(u,v)^2} \\
\lesssim \quad & \min_{s:V\to\mathbb{R}_{>0}} \alpha(G)^2 \frac{\sum_{(u,v)\in E} \pi(u)P(u,v)\rho_s(u,v)^2}{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)^2} \qquad \text{(Lemma 4.17)} \\
= \quad & \alpha(G)^2 \cdot \min_{s:V\to\mathbb{R}_{>0}} \frac{\sum_{(u,v)\in E} \pi(u)P(u,v)(s(u)+s(v))^2}{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)^2} \\
\leq \quad & \alpha(G)^2 \cdot \min_{s:V\to\mathbb{R}_{>0}} \frac{2\sum_{(u,v)\in E} \pi(u)P(u,v)(s(u)^2+s(v)^2)}{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)^2} \\
= \quad & \alpha(G)^2 \cdot \min_{s:V\to\mathbb{R}_{>0}} \frac{2\sum_{u} \pi(u)(\sum_v P(u,v))s(u)^2}{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)^2} \qquad (\pi\text{-reversibility of } P) \\
= \quad & \alpha(G)^2 \cdot \min_{s:V\to\mathbb{R}_{>0}} \frac{2\sum_{u} \pi(u)s(u)^2}{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)^2} \qquad (\text{Since } \sum_v P(u,v) = 1) \\
\leq \quad & \alpha(G)^2 \cdot \min_{s:V\to\mathbb{R}_{>0}} \frac{2\sum_{u} \pi(u)s(u)^2}{(\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v))^2} \qquad \text{(Cauchy-Schwarz)} \\
\lesssim \quad & \alpha(G)^2 \cdot \left[\max_{s:V\to\mathbb{R}_{>0}} \frac{\sum_{u,v} \pi(u)\pi(v)\rho_s(u,v)}{\sqrt{\sum_u \pi(u)s(u)^2}} \right]^{-2}.
\end{aligned}$$

□

Proof of Proposition 4.20 (Rounding to rational weight function). Let $M \in \mathbb{Z}$ be such that $M\pi(u) \geq 1/2$ for all $u \in V$. Then, set

$$\pi'(u) := \lceil M\pi(u) \rceil / M.$$

It is easy to see that $\pi'(u) \in \mathbb{Q}_{>0}$ and $\pi'(u) \geq \pi(u)$ for all $u \in V$. Moreover, since $M\pi(u) \geq 1/2$ we have $\lceil M\pi(u) \rceil \leq 2M\pi(u)$, and so $\pi'(u) \leq 2\pi(u)$ for all $u \in V$. □

Proof of Lemma 4.22 (Spike graph congestion). First, we prove the easier side, that there exists $B > 0$ such that

$$\min_{F_z} con_2(F_z) \leq B \min_{F'} con_2(F').$$

In fact, we can take $B = 1$. The idea is simple: from a solution F' to the MFP with demand graph $K_{n'}$ on the spike graph G', we construct a solution F_z to the MFP with demand graph $(K_n, \pi_z \times \pi_z)$ on G, with smaller congestion.

For any flow path p from u_i to v_j carrying some amount of flow, if $u = v$ we cancel the flow, else we shorten the two ends of the flow path (if needed) to make it go from u_0 to v_0. We do this to all flow paths in F', obtaining F_z. Refer to Figure 4 for an example. Observe that the operations only decrease the congestion of F', and in the end all flow paths go from u_0 to v_0 for some $u, v \in V$. By the canonical identification $u_0 \in V' \leftrightarrow u \in V$, we end up with a flow solution on G. Moreover, the amount of flow between u and $v \neq u$ in F is exactly $\pi_z(u) \times \pi_z(v)$. Therefore, the flow F_z satisfies the demand $(K_n, \pi_z \times \pi_z)$ and its congestion is no larger than that of F'. Hence,

$$\min_{F_z} con_2(F_z) \leq \min_{F'} con_2(F').$$

Figure 4: An illustration of the flow trimming procedure. For example, the flow path $c_1 \to c_0 \to b_0 \to b_1$ in purple gets trimmed to $c_0 \to b_0$.

Now, we prove the more difficult side, that there exists $A > 0$ such that

$$A \min_{F'} con_2(F') \le \min_{F_z} con_2(F_z).$$

We take $A = 1/6$. The proof proceeds by showing that the flow trimming procedure described above does not decrease $con_2(F')$ by too much. Let $c_{F'}(u_i)$ denote the congestion of vertex u_i with respect to flow solution F'.

First, we show that the congestion at u_0 dominates that at the other vertices u_i for $i \neq 0$:

$$\sum_{i \neq 0} c_{F'}(u_i) \le 2c_{F'}(u_0).$$

This is because, for any $i \neq 0$, (i) any flow path that passes through u_i must also pass through u_0 and (ii) in an optimal solution, any flow path that does not have u_i as an endpoint should not pass through u_i. Therefore, for any flow path, the contribution to u_0 is at least half of the contribution to the rest of u_i's, with equality case occurring when the flow path goes from u_i to u_j for some other $j \notin \{0, i\}$.

Therefore, if we let F'' be the flow solution obtained from F' by shortening flow paths $u_i \to v_j$ to $u_0 \to v_0$ (note that the flow paths $u_i \to u_j$ get shortened to $u_0 \to u_0$, which adds congestion to u_0), we have

$$9con_2(F'')^2 = \sum_{u \in V} (3c_{F'}(u_0))^2 \ge \sum_{u \in V} \left(\sum_{i < \pi_z(u)} c_{F'}(u_i) \right)^2 \ge \sum_{u_i \in V'} c_{F'}(u_i)^2 = con_2(F')^2.$$

Next, we show that the flow paths $u_0 \to u_0$ in F'' can all be discarded without decreasing the 2-congestion of the flow solution by too much. The total amount of congestion at u_0 due to the paths $u_0 \to u_0$ is

$$\sum_{0 \le i < j < \pi_z(u)} 1 = \frac{\pi_z(u)(\pi_z(u) - 1)}{2} < \pi_z(u)^2.$$

We only need to show that this is upper bounded by the total amount of congestion at u_0 due to other paths. Here we (finally) use the assumption that $\pi_{\max} < 1/2$. Check that this implies

$$\pi_z(u) \le \frac{1}{2} \sum_{v \in V} \pi_z(v)$$

for all $u \in V$. The congestion at u_0 due to other paths is at least the total amount of flow in F'' from u_0 to other vertices v_0, which is at least

$$\pi_z(u) \left(\sum_{v \in V \setminus \{u\}} \pi_z(v) \right) \ge \pi_z(u)^2.$$

Therefore, if we let F_z to be the flow solution obtained by dropping all $u_0 \to u_0$ flow paths from F'', we obtain

$$4con_2(F_z)^2 = \sum_{u \in V} (2c_{F_z}(u_0))^2 \geq \sum_{u \in V} (c_{F''}(u_0))^2 = con_2(F'')^2.$$

To sum up, the flow trimming procedure, which starts with a flow solution F' to the MFP with demand graph $K_{n'}$ on the spike graph G' and ends with a flow solution F_z to the MFP with demand graph $(K_n, \pi_z \times \pi_z)$ on G, satisfies

$$con_2(F_z) \geq \frac{1}{2} con_2(F'') \geq \frac{1}{6} con_2(F').$$

Since all flow solutions F_z to $(K_n, \pi_z \times \pi_z)$ can be obtained by applying such procedure to a flow solution F' to $K_{n'}$, taking minimum on both sides we have

$$\min_{F_z} con_2(F_z) \geq \frac{1}{6} \min_{F'} con_2(F'),$$

and the proof is complete. □

Proof of Lemma 4.23 (Spike graph topology). For the first part, given a genus g graph $G = (V, E)$, we can embed G in a genus g surface. We can then embed the spike graph G' in the same surface by attaching degree 1 vertices to G.

For the second part, we use contrapositive argument. Suppose the spike graph G' contains K_h as a minor, for some $h \geq 3$. Then, we can obtain K_h from G' by some sequence of edge deletion, vertex deletion, and edge contraction. The claim is that all vertices u_i, $i \geq 1$, eventually gets deleted. This is because contraction of edge (u_0, u_i) is the same as deleting u_i, deletion of edge (u_0, u_i) must be followed by deleting u_i (otherwise it becomes an isolated vertex), and K_h cannot contain any degree 1 vertex as $h \geq 3$. Furthermore, deletion of vertices u_i, $i \geq 1$, can be done in the very beginning of the sequence of operations.

Therefore, applying the sequence of operations on G', at some point we obtain G, then we end up with K_h. This implies that G contains K_h as a minor. □

References

[AHK10] Sanjeev Arora, Elad Hazan, and Satyen Kale. “$O(\log n)$ Approximation to SPARSEST CUT in $\tilde{O}(n^2)$ Time”. In: SIAM Journal on Computing 39.5 (2010), pp. 1748–1771.

[Alo86] Noga Alon. “Eigenvalues and expanders”. In: Combinatorica 6.2 (1986), pp. 83–96.

[AM85] Noga Alon and Vitali D Milman. “λ, isoperimetric inequalities for graphs, and superconcentrators”. In: Journal of Combinatorial Theory, Series B 38.1 (1985), pp. 73–88.

[ARV09] Sanjeev Arora, Satish Rao, and Umesh Vazirani. “Expander flows, geometric embeddings and graph partitioning”. In: Journal of the ACM (JACM) 56.2 (2009), pp. 1–37.

[Bar96] Yair Bartal. “Probabilistic approximation of metric spaces and its algorithmic applications”. In: Proceedings of 37th Conference on Foundations of Computer Science. IEEE. 1996, pp. 184–193.

[BLR10] Punyashloka Biswal, James R Lee, and Satish Rao. “Eigenvalue bounds, spectral partitioning, and metrical deformations via flows”. In: Journal of the ACM (JACM) 57.3 (2010), pp. 1–23.

[Che+22] Li Chen et al. “Maximum flow and minimum-cost flow in almost-linear time”. In: arXiv preprint arXiv:2203.00671 (2022).

[CS21] Julia Chuzhoy and Thatchaphol Saranurak. “Deterministic algorithms for decremental shortest paths via layered core decomposition”. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM. 2021, pp. 2478–2496.

[FHL08] Uriel Feige, MohammadTaghi Hajiaghayi, and James R Lee. “Improved approximation algorithms for minimum weight vertex separators”. In: SIAM Journal on Computing 38.2 (2008), pp. 629–657.

[JPV22] Vishesh Jain, Huy Tuan Pham, and Thuy-Duong Vuong. “Dimension reduction for maximum matchings and the Fastest Mixing Markov Chain”. In: arXiv preprint arXiv:2203.03858 (2022).

[Kel+11] Jonathan A Kelner et al. “Metric uniformization and spectral bounds for graphs”. In: Geometric and Functional Analysis 21.5 (2011), pp. 1117–1143.

[Kel+14] Jonathan A Kelner et al. “An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations”. In: Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms. SIAM. 2014, pp. 217–226.

[Kel06] Jonathan A Kelner. “Spectral partitioning, eigenvalue bounds, and circle packings for graphs of bounded genus”. In: SIAM Journal on Computing 35.4 (2006), pp. 882–902.

[KLT22] Tsz Chiu Kwok, Lap Chi Lau, and Kam Chuen Tung. “Cheeger Inequalities for Vertex Expansion and Reweighted Eigenvalues”. In: arXiv preprint arXiv:2203.06168 (2022).

[KPR93] Philip Klein, Serge A Plotkin, and Satish Rao. “Excluded minors, network decomposition, and multicommodity flow”. In: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing. 1993, pp. 682–690.

[KRV09] Rohit Khandekar, Satish Rao, and Umesh Vazirani. “Graph partitioning using single commodity flows”. In: Journal of the ACM (JACM) 56.4 (2009), pp. 1–15.

[LR99] Tom Leighton and Satish Rao. “Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms”. In: Journal of the ACM (JACM) 46.6 (1999), pp. 787–832.

[LS10] James R Lee and Anastasios Sidiropoulos. “Genus and the geometry of the cut graph”. In: Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms. SIAM. 2010, pp. 193–201.

[LT79] Richard J Lipton and Robert Endre Tarjan. "A separator theorem for planar graphs". In: SIAM Journal on Applied Mathematics 36.2 (1979), pp. 177–189.

[Ore+08] Lorenzo Orecchia et al. "On partitioning graphs via single commodity flows". In: Proceedings of the fortieth annual ACM symposium on Theory of computing. 2008, pp. 461–470.

[OZ22] Sam Olesker-Taylor and Luca Zanetti. "Geometric Bounds on the Fastest Mixing Markov Chain". In: 13th Innovations in Theoretical Computer Science Conference (ITCS 2022). Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2022.

[Roc05] Sébastien Roch. "Bounding fastest mixing". In: Electronic Communications in Probability 10 (2005), pp. 282–296.

[She09] Jonah Sherman. "Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$-approximations to sparsest cut". In: 2009 50th Annual IEEE Symposium on Foundations of Computer Science. IEEE. 2009, pp. 363–372.

[ST96] Daniel A Spielman and Shang-Hua Teng. "Spectral partitioning works: Planar graphs and finite element meshes". In: Proceedings of 37th Conference on Foundations of Computer Science. IEEE. 1996, pp. 96–105.

[Tan84] R Michael Tanner. "Explicit concentrators from generalized N-gons". In: SIAM Journal on Algebraic Discrete Methods 5.3 (1984), pp. 287–293.