Discrete Mathematics 349 (2026) 115298

Contents lists available at ScienceDirect

Discrete Mathematics

journal homepage: www.elsevier.com/locate/disc

A Cheeger inequality for size-specific conductance ☆

Yufan Huang *, David F. Gleich *

Department of Computer Science, Purdue University, 305 N University St, West Lafayette, 47907, IN, United States

A R T I C L E   I N F O

Article history:
Received 12 March 2025
Received in revised form 29 January 2026
Accepted 19 June 2026
Available online xxxx

Keywords:
Graph theory
Isoperimetric inequality

A B S T R A C T

The μ-conductance measure proposed by Lovász and Simonovits is a size-specific conductance score that identifies the set with smallest conductance while disregarding those sets with volume smaller than a μ fraction of the whole graph. Using μ-conductance enables us to study the network structures in new ways. In this manuscript we study a modified spectral cut for μ-conductance that is a natural relaxation of the integer program of μ-conductance and show that the optimum of this program has a two-sided Cheeger inequality with μ-conductance.

1. Introduction

Graph clustering is one of the most fundamental problems in graph analysis and has many practical applications [30]. A major class of clustering methods is based on finding small cuts, motivated by the observation that meaningful clusters tend to be internally well-connected while being sparsely connected to the rest of the graph. The minimum cut problem is perhaps the most classical example, but it often produces highly unbalanced solutions, such as isolating a single low-degree vertex [16]. To address this issue, the sparsest cut (or low-conductance cut) objective instead minimizes the ratio between the size of the cut, $|\partial S|$, and the size of the smaller side of the partition, $|S|$ (or, in the conductance case, the smaller volume side), and has achieved notable success [13,25].

However, a limitation of the sparsest cut and low-conductance cut objectives is that they capture only a single structure in a graph—the set with the worst bottleneck. In reality, the relationship between cut size $|\partial S|$ and set size $|S|$ can be much more complex. Leskovec et al. [26] conducted a systematic study of the so-called *network community profile*, which measures the minimum conductance over a range of different set sizes, across a large collection of real-world graphs. They observed that, in many graphs, the minimum-conductance sets at different size scales can have dramatically different conductance values. For example, the minimum conductance set of size roughly half of the graph can have conductance several orders of magnitude larger than that of the unconstrained minimum-conductance set.

A related size-specific cut measure is μ-conductance, originally introduced by Lovász and Simonovits in the study of Markov chains for sampling convex bodies [28]. μ-conductance seeks the minimum-conductance set subject to a size constraint: the volume of the set S, denoted $\mathrm{Vol}(S)$, must be at least a μ fraction of the total graph volume. By varying μ, one can reveal richer information about the cut structure of a graph while ignoring sets of extremely small volume.

Despite their descriptive power, such size-specific cut profiles cannot be computed exactly, as computing minimum-conductance sets is NP-hard [31]. A common heuristic approach is to generate many graph cuts using different low-

☆ The work was funded in part by NSF CCF-1909528, IIS-2007481, DOE DE-SC0023162, and the IARPA Agile program. We thank C. Seshadhri for discussions on this topic.

* Corresponding authors.
E-mail addresses: 2019hyf@gmail.com (Y. Huang), dgleich@purdue.edu (D.F. Gleich).

https://doi.org/10.1016/j.disc.2026.115298

conductance algorithms and aggregate the results [26]. This approach faces two main challenges. First, existing approximation algorithms for low-conductance sets often provide guarantees only in asymptotic Big-O form, for example $O(\log n)$ or $O(\sqrt{\log n})$, which obscures the quality of the resulting cuts in practice [3,20,25]. Second, algorithms for finding low-conductance sets under size or volume constraints are typically bicriteria, meaning they do not enforce the size constraint exactly [2–4]. As a result, existing size-specific cut profiles serve only as proxies (upper bounds) for the true profiles. Huang et al. [19] addressed this issue by deriving a lower bound via a relaxation of the integer program for μ-conductance and demonstrated that this bound can be computed on large real-world graphs with millions of nodes and tens of millions of edges.

In this paper, we show that the spectral program studied in Huang et al. [19] satisfies a Cheeger-type inequality with μ-conductance. Cheeger's inequality is a foundational result in spectral graph theory [11], relating graph conductance to the eigenvalues of the normalized Laplacian. The original inequality of +Cheeger [11] bounds the smallest Laplacian eigenvalue of a +Riemannian manifold; its discrete analogue for graphs appeared later, +with the earliest instances due to Dodziuk [15] and +Alon and Milman [1]. It has played a central role in the design of graph algorithms and has been extended in many directions, including directed graphs and hypergraphs [6,23], multiway cuts [27], higher-order spectral gaps [22,24], general types of diffusions [17], and conductance across multiple graphs [21] (which can be specialized for useful MinCut-specific bounds [32]).

The classical Cheeger inequality states that

$$2\phi_G \geq \lambda_G \geq \phi_G^2/2,$$

where λ_G is the spectral gap of the normalized Laplacian and ϕ_G is the minimum conductance over all vertex sets in G.

Let λ_μ denote the optimum of the spectral program for μ-conductance and let ϕ_μ denote the μ-conductance. We prove that

$$2\phi_\mu \geq \lambda_\mu \geq \frac{1}{4}\left(\frac{\mu^2\phi_\mu + (1/2-\mu)\,\phi_G}{\mu^2 + 1/2 - \mu}\right)^2.$$

This result can be viewed as a size-specific generalization of the classical Cheeger inequality, relating μ-conductance to the optimum value of the corresponding spectral relaxation. When $\mu = 0$, the size constraint is vacuous, and we recover the standard setting: $\phi_\mu = \phi_G$, $\lambda_\mu = \lambda_G$, and the inequality reduces to the classical Cheeger bound. In contrast, as μ approaches $1/2$, corresponding to increasingly balanced cuts, our bound implies

$$2\phi_\mu \geq \lambda_\mu \geq \phi_\mu^2/4.$$

2. Preliminaries

In a weighted undirected graph $G = (V, E, w : E \to \mathbb{R})$ with n vertices and m edges, each edge $uv \in E$ has weight $w(u, v) \in \mathbb{R}$. The (weighted) degree of a vertex v is defined as the sum of weights of its incident edges: $d(v) \overset{\text{def}}{=} \sum_{u\in V: uv\in E} w(v, u)$.

For a vertex set $S \subseteq V$, its *volume* is the total degree of vertices in S: $\mathrm{Vol}(S) \overset{\text{def}}{=} \sum_{v\in S} d(v)$. This serves as a size measure analogous to $|S|$. The volume of the entire graph is $\mathrm{Vol}(G) \overset{\text{def}}{=} \mathrm{Vol}(V) = \sum_{v\in V} d(v)$. The set of edges with one endpoint in S and the other in $\bar{S} = V \setminus S$ is the *cut* separating S from $\bar{S}$: $\partial S \overset{\text{def}}{=} \{uv \in E : u \in S, v \in \bar{S}\}$, with total weight $|\partial S| \overset{\text{def}}{=} \sum_{uv\in\partial S} w(u, v)$. Finding sparse cuts is central to community detection, since different communities are often weakly connected. However, the raw cut size ignores the relative sizes of S and $\bar{S}$, which can lead to highly unbalanced partitions. To address this, *conductance* normalizes cut size by the smaller volume:

$$\phi(S) \overset{\text{def}}{=} \frac{|\partial S|}{\min\{\mathrm{Vol}(S), \mathrm{Vol}(\bar{S})\}}.$$

Low-conductance sets tend to induce sparse and more balanced cuts, since $\mathrm{Vol}(S) + \mathrm{Vol}(\bar{S}) = \mathrm{Vol}(G)$. The conductance of the graph is then $\phi_G \overset{\text{def}}{=} \min_{S\subset V} \phi(S)$, which measures the overall connectivity of G.

An important toolkit for studying graph structure, for instance conductance, is *spectral graph theory*, which focuses on graph-associated matrices. The *degree matrix* of G is the diagonal matrix $\boldsymbol{D}$ with $\boldsymbol{D}_{vv} = d(v)$. The *graph Laplacian* is

$$\boldsymbol{L} \overset{\text{def}}{=} \sum_{uv\in E} w(u, v)\,(\mathbf{1}_u - \mathbf{1}_v)(\mathbf{1}_u - \mathbf{1}_v)^T,$$

where $\mathbf{1}_S$ denotes the indicator vector of a set S, and $\mathbf{1}_a = \mathbf{1}_{\{a\}}$. The normalized Laplacian is $\boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2}$.

3. Cheeger inequality for μ-conductance

In this section, we first define the μ-conductance, which provides a more fine-grained view of a graph's cut structure. In contrast, the standard notion of conductance captures only a single aspect of the graph—namely, the set corresponding to the worst bottleneck. We then introduce the modified spectral program inspired by the μ-conductance integer program from Huang et al. [19]. Finally, we show that the optimum of this spectral program satisfies a two-sided Cheeger inequality with respect to μ-conductance.

3.1. The definition of μ-conductance

The idea of μ-conductance is a parameterized variant of conductance that arises from Markov chain theory [28]. Basically, it computes the best set with the smallest conductance disregarding those sets with volume smaller than $\mu\mathrm{Vol}(G)$. Formally, it is defined as[1]

$$\begin{aligned} \phi_\mu(G) = \ & \underset{S\subset V}{\text{minimize}} \ \ \phi(S) \\ & \text{subject to} \ \ \mu\mathrm{Vol}(G) \le \mathrm{Vol}(S) \le (1-\mu)\mathrm{Vol}(G). \end{aligned} \tag{1}$$

In the notation of μ-conductance, the original conductance ϕ_G can be expressed as $\phi_0(G)$. So by computing μ-conductance for multiple μs, we may be able to gain additional information about a graph. This will only be productive if ϕ_G arises *only* from a small set in the graph. If $\phi(G)$ arises from a set of nearly $\mathrm{Vol}(G)/2$, then there are no differences between conductance and μ-conductance.

3.2. A spectral program for μ-conductance

Before we introduce our modified spectral program for μ-conductance, let us first revisit the relationship between conductance and the spectral cut to smooth the transition to our new program. Basically, the problem of finding the set of smallest conductance is equivalent to (up to a constant factor) the following integer program

$$\begin{aligned} & \underset{S\subset V}{\text{minimize}} \quad \frac{\boldsymbol{\psi}^T \boldsymbol{L}\boldsymbol{\psi}}{\boldsymbol{\psi}^T \boldsymbol{D}\boldsymbol{\psi}} \\ & \text{subject to} \quad \boldsymbol{\psi} = \frac{1}{\mathrm{Vol}(S)}\mathbf{1}_S - \frac{1}{\mathrm{Vol}(\bar{S})}\mathbf{1}_{\bar{S}} \end{aligned} \tag{2}$$

where $\boldsymbol{\psi}$ is a shifted and scaled indicator vector for the vertex set S. The spectral cut

$$\begin{aligned} & \underset{\mathbf{x}\in\mathbb{R}^{|V|}}{\text{minimize}} \quad \frac{\mathbf{x}^T \boldsymbol{L}\mathbf{x}}{\mathbf{x}^T \boldsymbol{D}\mathbf{x}} \\ & \text{subject to} \quad \mathbf{x}^T\mathbf{d} = 0, \end{aligned} \tag{3}$$

is actually a relaxation of (2) by replacing $\boldsymbol{\psi}$ with vectors orthogonal to $\mathbf{d}$.

Notice that the objective in (2) is scale-invariant with regard to $\boldsymbol{\psi}$, it can be rewritten as

$$\begin{aligned} & \underset{S\subset V}{\text{minimize}} \quad \boldsymbol{\psi}^T \boldsymbol{L}\boldsymbol{\psi} \\ & \text{subject to} \quad \boldsymbol{\psi}^T \boldsymbol{D}\boldsymbol{\psi} = 1 \\ & \qquad\qquad\quad \boldsymbol{\psi} = \sqrt{\frac{\mathrm{Vol}(S)\mathrm{Vol}(\bar{S})}{\mathrm{Vol}(G)}}\left(\frac{1}{\mathrm{Vol}(S)}\mathbf{1}_S - \frac{1}{\mathrm{Vol}(\bar{S})}\mathbf{1}_{\bar{S}}\right), \end{aligned}$$

where $\boldsymbol{\psi}$ is normalized to satisfy $\boldsymbol{\psi}^T\boldsymbol{D}\boldsymbol{\psi} = 1$.

[1] Here we adopt a slightly different definition from the original paper. The original μ-conductance is defined as

$$\phi_\mu(G) = \min_{S:\mathrm{Vol}(S)\ge\mu, \mathrm{Vol}(\bar{S})\ge\mu} |\partial S| \,/\, \min\left\{\mathrm{Vol}(S)-\mu, \mathrm{Vol}(\bar{S})-\mu\right\}.$$

They are similar in spirit as they both neglect sets with volume smaller than a specific volume but the original one involves a perturbed conductance. The one we adopt keeps the standard conductance definition and only adds the size constraint. We believe this provides a cleaner picture in network analysis.

We have a similar integer program for μ-conductance, simply adding additional volume constraints on S,

$$\begin{aligned} \underset{S\subset V}{\text{minimize}} \quad & \boldsymbol{\psi}^T \boldsymbol{L}\boldsymbol{\psi} \\ \text{subject to} \quad & \boldsymbol{\psi} = \sqrt{\frac{\text{Vol}(\bar{S})}{\text{Vol}(S)\text{Vol}(G)}}\mathbf{1}_S - \sqrt{\frac{\text{Vol}(S)}{\text{Vol}(\bar{S})\text{Vol}(G)}}\mathbf{1}_{\bar{S}} \\ & \boldsymbol{\psi}^T \boldsymbol{D}\boldsymbol{\psi} = 1 \\ & \mu\text{Vol}(G) \le \text{Vol}(S) \le (1-\mu)\text{Vol}(G). \end{aligned} \tag{4}$$

Notice that the constraint $\mu\text{Vol}(G) \le \text{Vol}(S) \le (1-\mu)\text{Vol}(G)$ implies that the entries of $\boldsymbol{\psi}$ must be delocalized, in other words there will be no entries with very large magnitude and no entries with very small magnitude. Hence this integer program can be relaxed to the following spectral program

$$\begin{aligned} \lambda_\mu = \underset{\mathbf{x}\in\mathbb{R}^{|V|}}{\text{minimize}} \quad & \mathbf{x}^T \boldsymbol{L}\mathbf{x} \\ \text{subject to} \quad & \mathbf{x}^T \mathbf{d} = 0 \\ & \mathbf{x}^T \boldsymbol{D}\mathbf{x} = 1 \\ & \|\mathbf{x}\|_\infty \le \sqrt{\frac{1-\mu}{\mu\text{Vol}(G)}} \\ & |\mathbf{x}_i| \ge \sqrt{\frac{\mu}{(1-\mu)\text{Vol}(G)}}, \quad \forall i \in V. \end{aligned} \tag{5}$$

3.3. Main result

In this section, we demonstrate that a non-trivial relation is preserved by the relaxation from (4) to (5) and the optimum of spectral program (5) has a two-sided Cheeger inequality with μ-conductance. Our main result is the following Theorem.

Theorem 3.1. *Given a graph G and a constant $0 \le \mu \le 1/2$, we have*

$$2\phi_\mu \ge \lambda_\mu \ge \frac{1}{4}\left(\frac{\mu^2\phi_\mu + (1/2-\mu)\,\phi_0}{\mu^2 + 1/2 - \mu}\right)^2$$

For simplicity, we assume that there exists a set S with volume between $\mu\text{Vol}(G)$ and $(1-\mu)\text{Vol}(G)$; otherwise, the notion of μ-conductance for this value of μ is not meaningful. The lower bound is essentially the square of an interpolation between ϕ_μ and ϕ_0.

At the two extreme cases, when $\mu = 0$ we have $2\phi_0 \ge \lambda_0 \ge \phi_0^2/4$ (the standard Cheeger up to a constant factor), and when $\mu = 1/2$ we have $2\phi_{1/2} \ge \lambda_{1/2} \ge \phi_{1/2}^2/4$.

We point out one disadvantage of our result. Unlike the linear upper bound $\lambda_\mu \le 2\phi_\mu$, our bound does not yield a square-root upper bound on ϕ_μ in terms of λ_μ that holds uniformly in μ, as the coefficient $\mu^2/(\mu^2+1/2-\mu)$ of ϕ_μ in the lower bound goes to 0 as $\mu \downarrow 0$.

We split the proof of Theorem 3.1 into two parts, the lower bound of λ_μ and the upper bound of λ_μ.

3.4. Proof for the upper bound of λ_μ

This is the easier side of Theorem 3.1, of which the proof was also included in Huang et al. [19]. We include it here for completeness. The proof directly follows from the fact that (5) is one relaxation of (4).

Lemma 3.2. *Given a graph G and a constant $0 \le \mu \le 1/2$, we have*

$$2\phi_\mu \ge \lambda_\mu.$$

Proof. Let S be the vertex set which achieves the optimal μ-conductance. By our relaxation, we know

$$\boldsymbol{\psi} = \sqrt{\frac{\text{Vol}(\bar{S})}{\text{Vol}(S)\text{Vol}(G)}}\mathbf{1}_S - \sqrt{\frac{\text{Vol}(S)}{\text{Vol}(\bar{S})\text{Vol}(G)}}\mathbf{1}_{\bar{S}}$$

is naturally in the feasible region of (5). Then, the corresponding objective value is

$$\begin{aligned}\boldsymbol{\psi}^T \boldsymbol{L}\boldsymbol{\psi} &= \frac{|\partial S|\mathrm{Vol}(G)}{\mathrm{Vol}(S)\mathrm{Vol}(\bar{S})} \\ &\le \frac{2|\partial S|}{\min\{\mathrm{Vol}(S), \mathrm{Vol}(\bar{S})\}} \\ &= 2\phi_\mu,\end{aligned}$$

which implies $\lambda_\mu \le 2\phi_\mu$. □

3.5. Proof for lower bound of λ_μ

Now we turn to the challenging part of the proof, lower bounding λ_μ via ϕ_μ. Our proof is mainly inspired by the proof in Chung [12], with a few significant differences that we will point out in the proof.

Let $\mathbf{g}$ denote the vector that achieves the optimum of (5). We order all vertices such that $\mathbf{g}(v_1) \ge \mathbf{g}(v_2) \ge \cdots \ge \mathbf{g}(v_n)$ and let $S_j = \{v_1, v_2, \ldots, v_j\}$ denote the set of first j vertices for $j \in [n]$. Let h be the largest index such that $\mathrm{Vol}(S_h) \le \mathrm{Vol}(G)/2$ (this implies that $h < n$). As a result, we have

$$\min\left\{\mathrm{Vol}(S_i), \mathrm{Vol}(\bar{S}_i)\right\} = \begin{cases}\mathrm{Vol}(S_i), & \forall i \le h \\ \mathrm{Vol}(\bar{S}_i), & \forall i > h\end{cases}. \tag{6}$$

Then, we decompose $\mathbf{g}$ into two vectors, denoted by $\mathbf{g}_+$ and $\mathbf{g}_-$ where

$$\begin{aligned}\mathbf{g}_+(v_i) &= \max(\mathbf{g}(v_i) - \mathbf{g}(v_{h+1}), 0), \\ \mathbf{g}_-(v_i) &= \max(\mathbf{g}(v_{h+1}) - \mathbf{g}(v_i), 0).\end{aligned}$$

Intuitively, $\mathbf{g}_+$ and $\mathbf{g}_-$ correspond to the positive and negative components of "shifted" $\mathbf{g}$ with $\mathbf{g}(v_{h+1})$ being the new zero point. By definition, we have

$$\mathbf{g}_+(v_1) \ge \ldots \ge \mathbf{g}_+(v_h) \ge \mathbf{g}_+(v_{h+1}) = \ldots = \mathbf{g}_+(v_n) = 0$$

and

$$\mathbf{g}_-(v_n) \ge \ldots \ge \mathbf{g}_-(v_{h+2}) \ge \mathbf{g}_-(v_{h+1}) = \ldots = \mathbf{g}_-(v_1) = 0.$$

For any $\mathbf{x} \in \mathbb{R}^V$, let

$$R(\mathbf{x}) \stackrel{\text{def}}{=} \frac{\mathbf{x}^T \boldsymbol{L}\mathbf{x}}{\mathbf{x}^T \boldsymbol{D}\mathbf{x}}.$$

On one hand, because $\sum_v \mathbf{g}(v)d(v) = 0$,

$$\mathbf{g}(v_{h+1})\sum_v \mathbf{g}(v)d(v) = 0,$$

thus

$$\begin{aligned}&\sum_v \left(\mathbf{g}_+(v)^2 d(v) + \mathbf{g}_-(v)^2 d(v)\right) \\ &= \sum_v \left(\mathbf{g}(v) - \mathbf{g}(v_{h+1})\right)^2 d(v) \\ &= \sum_v \mathbf{g}(v)^2 d(v) - 2\sum_v \mathbf{g}(v)\mathbf{g}(v_{h+1})d(v) + \sum_v \mathbf{g}(v_{h+1})^2 d(v) \\ &\ge \sum_v \mathbf{g}(v)^2 d(v).\end{aligned}$$

On the other hand, with some casework, we have

$$(\mathbf{g}(u) - \mathbf{g}(v))^2 \le \left(\mathbf{g}_+(u) - \mathbf{g}_+(v)\right)^2 + \left(\mathbf{g}_-(u) - \mathbf{g}_-(v)\right)^2.$$

Together we have

$$\lambda_\mu = R(\mathbf{g}) = \frac{\sum_{uv\in E}(\mathbf{g}(u)-\mathbf{g}(v))^2}{\sum_{v\in V} d(v)\mathbf{g}(v)^2} \geq \frac{\sum_{uv\in E}(\mathbf{g}_+(u)-\mathbf{g}_+(v))^2 + \sum_{uv\in E}(\mathbf{g}_-(u)-\mathbf{g}_-(v))^2}{\sum_{v\in V} d(v)\mathbf{g}_+(v)^2 + \sum_{v\in V} d(v)\mathbf{g}_-(v)^2}. \tag{7}$$

In Chung [12], to get $\lambda_0 \geq \phi_0^2/2$, one can analyze either $R(\mathbf{g}_+)$ or $R(\mathbf{g}_-)$ because (7) implies that $R(\mathbf{g}) \geq \min\{R(\mathbf{g}_+), R(\mathbf{g}_-)\}$ and $R(\mathbf{g}_+), R(\mathbf{g}_-)$ are symmetric. However, simply considering one of $R(\mathbf{g}_+)$ and $R(\mathbf{g}_-)$ can not give a clean lower bound involving ϕ_μ because the entries of $\mathbf{g}_+$ or $\mathbf{g}_-$ can be arranged in a way that $R(\mathbf{g}_+)$ or $R(\mathbf{g}_-)$ is overly relaxed, which gives the trivial lower bound $\phi_0^2/2$. So in our proof, we take $\mathbf{g}_+$ and $\mathbf{g}_-$ into account simultaneously. In this way the entries of $\mathbf{g}_+$ and $\mathbf{g}_-$ have less freedom and we get a lower bound tighter than $\phi_0^2/2$, in particular with ϕ_μ involved.

For convenience, we let $\mathbf{p}_u$ denote $\mathbf{g}_+(u)$ and $\mathbf{q}_u$ denote $\mathbf{g}_-(u)$. When there is no risk of confusion, we abuse the notations to let $\mathbf{p}_i$ be $\mathbf{p}_{v_i}$ and $\mathbf{q}_i$ be $\mathbf{q}_{v_i}$. As a result, (7) can be written as

$$\lambda_\mu \geq \frac{\sum_{uv\in E}\left(\mathbf{p}_u-\mathbf{p}_v\right)^2 + \sum_{uv\in E}\left(\mathbf{q}_u-\mathbf{q}_v\right)^2}{\sum_{v\in V} d(v)\mathbf{p}_u^2 + \sum_{v\in V} d(v)\mathbf{q}_u^2} \tag{8}$$

We multiply the numerator and denominator of (7) both by

$$\sum_{uv\in E}\left(\mathbf{p}_u+\mathbf{p}_v\right)^2 + \left(\mathbf{q}_u+\mathbf{q}_v\right)^2.$$

By Cauchy–Schwarz inequality, the numerator is lower bounded by

$$\begin{aligned}
&\sum_{uv\in E}\left(\left(\mathbf{p}_u-\mathbf{p}_v\right)^2+\left(\mathbf{q}_u-\mathbf{q}_v\right)^2\right)\sum_{uv\in E}\left(\left(\mathbf{p}_u+\mathbf{p}_v\right)^2+\left(\mathbf{q}_u+\mathbf{q}_v\right)^2\right)\\
&\geq \left(\sum_{uv\in E}\sqrt{\left(\mathbf{p}_u-\mathbf{p}_v\right)^2+\left(\mathbf{q}_u-\mathbf{q}_v\right)^2}\sqrt{\left(\mathbf{p}_u+\mathbf{p}_v\right)^2+\left(\mathbf{q}_u+\mathbf{q}_v\right)^2}\right)^2\\
&\overset{(a)}{\geq} \left(\sum_{uv\in E}\sqrt{\left(\mathbf{p}_u^2-\mathbf{p}_v^2\right)^2+\left(\mathbf{q}_u^2-\mathbf{q}_v^2\right)^2}\right)^2\\
&\overset{(b)}{\geq} \left(\sum_{uv\in E}\frac{\left|\mathbf{p}_u^2-\mathbf{p}_v^2\right|+\left|\mathbf{q}_u^2-\mathbf{q}_v^2\right|}{\sqrt{2}}\right)^2\\
&\overset{(c)}{=} \frac{1}{2}\left(\sum_{i=1}^{n-1}\left(\mathbf{p}_i^2-\mathbf{p}_{i+1}^2\right)|\partial S_i| + \sum_{i=1}^{n-1}\left(\mathbf{q}_{i+1}^2-\mathbf{q}_i^2\right)|\partial S_i|\right)^2
\end{aligned} \tag{9}$$

where (a) is by expanding the product and dropping nonnegative cross-terms, (b) is because of AM-QM inequality, and (c) is from rearranging terms.

Similarly, the denominator is upper bounded by

$$\begin{aligned}
&\left(\sum_{v\in V} d(v)\mathbf{p}_v^2 + \sum_{v\in V} d(v)\mathbf{q}_v^2\right)\left(\sum_{uv\in E}\left(\left(\mathbf{p}_u+\mathbf{p}_v\right)^2+\left(\mathbf{q}_u+\mathbf{q}_v\right)^2\right)\right)\\
&\overset{(d)}{\leq} \left(\sum_{v\in V} d(v)\mathbf{p}_v^2 + \sum_{v\in V} d(v)\mathbf{q}_v^2\right)\left(\sum_{uv\in E} 2\left(\mathbf{p}_u^2+\mathbf{p}_v^2\right)+2\left(\mathbf{q}_u^2+\mathbf{q}_v^2\right)\right)\\
&\overset{(e)}{=} 2\left(\sum_{v\in V} d(v)\left(\mathbf{p}_v^2+\mathbf{q}_v^2\right)\right)^2
\end{aligned} \tag{10}$$

where (d) is by AM-QM inequality and (e) is by rearranging terms.

Plugging (9) and (10) back into (8) gives

$$\lambda_\mu \geq \frac{1}{4}\left(\frac{\sum_{i=1}^{n-1}\left(\mathbf{p}_i^2-\mathbf{p}_{i+1}^2\right)|\partial S_i| + \sum_{i=1}^{n-1}\left(\mathbf{q}_{i+1}^2-\mathbf{q}_i^2\right)|\partial S_i|}{\sum_{v\in V} d(v)\left(\mathbf{p}_v^2+\mathbf{q}_v^2\right)}\right)^2. \tag{11}$$

Now our goal is to lower bound (11).

Table 1
Three key indices in our proof.

Index	Definition	Interpretation
h	largest index s. t. $\mathrm{Vol}(S_h) \le \mathrm{Vol}(G)/2$	*half*
x	largest index s. t. $\mathrm{Vol}(S_x) < \mu\mathrm{Vol}(G)$	μ fraction
y	largest index s. t. $\mathrm{Vol}(S_y) \le (1-\mu)\mathrm{Vol}(G)$	$(1-\mu)$ fraction

Before doing that, let us introduce a few indices which will repeatedly show up in the analysis. Besides h, which is the largest integer such that $\mathrm{Vol}(S_h) \le \mathrm{Vol}(G)/2$ as introduced before, we let x, y be the largest indices such that $\mathrm{Vol}(S_x) < \mu\mathrm{Vol}(G), \mathrm{Vol}(S_y) \le (1-\mu)\mathrm{Vol}(G)$, respectively. We need these two indices because this time we cannot loosely lower bound every $|\partial S_i|$ by $\phi_0 \min\left\{\mathrm{Vol}(S_i), \mathrm{Vol}(\bar{S}_i)\right\}$, we need a tighter bound involving ϕ_μ. Therefore we need to do some casework depending on $\mathrm{Vol}(S_i)$, where x and y are two important cutoffs. These three indices are summarized in Table 1. Furthermore, without loss of generality, we assume $\mathbf{g}(v_{h+1}) < 0$; the case $\mathbf{g}(v_{h+1}) \ge 0$ is symmetric because replacing $\mathbf{g}$ with $-\mathbf{g}$ gives the same objective.

With the constraints on the magnitudes of the entries and the condition $\mathbf{g}^T\mathbf{d} = 0$, intuitively there cannot be too many negative entries. In fact, $\mathbf{g}(v_{x+1})$ is nonnegative and we have the following simple fact on its magnitude.

Lemma 3.3. *By definition of x, we have*

$$\mathbf{g}(v_{x+1}) \ge \sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}.$$

Proof. Assume instead $\mathbf{g}(v_{x+1}) < 0$, then

$$\begin{aligned}\sum_v d(v)\mathbf{g}(v) &= \sum_{i=1}^{x} d(v)\mathbf{g}(v_i) + \sum_{i=x+1}^{n} d(v_i)\mathbf{g}(v_i) \\ &\overset{(a)}{\le} \mathrm{Vol}(S_x)\mathbf{g}(v_1) + (\mathrm{Vol}(G) - \mathrm{Vol}(S_x))\mathbf{g}(v_{x+1}) \\ &\overset{(b)}{\le} \mathrm{Vol}(S_x)\sqrt{\frac{1-\mu}{\mu\mathrm{Vol}(G)}} - (\mathrm{Vol}(G) - \mathrm{Vol}(S_x))\sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}} \\ &\overset{(c)}{<} \sqrt{\mu(1-\mu)\mathrm{Vol}(G)} - \sqrt{\mu(1-\mu)\mathrm{Vol}(G)} = 0\end{aligned}$$

where (a) is because the entries of $\mathbf{g}$ are sorted in the decreasing order, (b) is due to the constraints in (5) on the magnitude of the entries in $\mathbf{g}$, and (c) follows from the definition of x. This contradicts the constraint $\mathbf{g}^T\mathbf{d} = 0$. Therefore we know $\mathbf{g}(v_{x+1}) \ge 0$ and by the lower bound on $|\mathbf{g}_i|$ in (5), we get the desired property. □

We now show that the ratio in (11) is bounded below by

$$\frac{\sum_{i=1}^{n-1}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)|\partial S_i| + \sum_{i=1}^{n-1}\left(\mathbf{q}_{i+1}^2 - \mathbf{q}_i^2\right)|\partial S_i|}{\sum_{v\in V} d(v)\left(\mathbf{p}_v^2 + \mathbf{q}_v^2\right)} \ge \frac{\mu^2\phi_\mu + (1/2-\mu)\phi_0}{\mu^2 - \mu + 1/2}, \tag{12}$$

which, plugged into (11), proves the lower bound of λ_μ in Theorem 3.1. Before diving into the technical proof, let us gain some intuition about the coefficients of ϕ_μ, ϕ_0 in the bound. Consider a graph with all vertices having roughly the same and relatively small degrees, and a vector $\mathbf{g}$ with

$$\mathbf{g}(v_i) \approx \begin{cases} \sqrt{\frac{1-\mu}{\mu\mathrm{Vol}(G)}}, & 1 \le i \le x \\ \sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}, & i = x+1 \\ -\sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}, & i > x+1 \end{cases}.$$

After some slight value perturbation, $\mathbf{g}$ can satisfy $\mathbf{g}^T\mathbf{d} = 0, \mathbf{g}^T\boldsymbol{D}\mathbf{g} = 1$ because $\mathrm{Vol}(S_x) \approx \mathrm{Vol}(S_{x+1}) \approx \mu\mathrm{Vol}(G)$. As a result, $\mathbf{g}$ is a feasible solution. And we have

$$\mathbf{p}_i \approx \begin{cases} \sqrt{\frac{1-\mu}{\mu\mathrm{Vol}(G)}} + \sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}, & 1 \le i \le x \\ 2\sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}, & i = x+1 \\ 0, & i > x+1 \end{cases},$$

and $\mathbf{q}_i = 0$ for any i.

On one hand, for this feasible $\mathbf{g}$, the numerator of the left side in (12) becomes roughly $(\mathbf{p}_x^2 - \mathbf{p}_{x+1}^2)|\partial S_x| + \mathbf{p}_{x+1}^2|\partial S_{x+1}|$ and is lower bounded by

$$(\mathbf{p}_x^2 - \mathbf{p}_{x+1}^2)\phi_0 \mathrm{Vol}(S_x) + \mathbf{p}_{x+1}^2 \phi_\mu \mathrm{Vol}(S_{x+1}) \approx \frac{1-4\mu^2}{1-\mu}\phi_0 + \frac{4\mu^2}{1-\mu}\phi_\mu.$$

On the other hand, with some computation, we see that the denominator is roughly $1/(1-\mu)$. Overall, the lower bound for this $\mathbf{g}$ is roughly $4\mu^2\phi_\mu + (1-4\mu^2)\phi_0$, and the lower bound is tight when $|\partial S_x| \approx \phi_0 \mathrm{Vol}(S_x)$, $|\partial S_{x+1}| \approx \phi_\mu \mathrm{Vol}(S_{x+1})$.

Although this does not rule out the possibility of having a tighter lower bound with a totally different proof technique, this implies that the lower bound we give is likely not possible to dramatically improve within the current proof framework. In other words, applying Cauchy–Schwarz and then considering threshold (sweeping) cuts $|\partial S_i|$ is unlikely to give a strong lower bound.

Now, with this intuition, we establish (12). Observe that

$$\begin{aligned}
&\sum_{i=1}^{n-1}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)|\partial S_i| \\
&\overset{(a)}{=} \sum_{i=1}^{x}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)|\partial S_i| + \sum_{i=x+1}^{h}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)|\partial S_i| \\
&\overset{(b)}{\geq} \phi_0 \sum_{i=1}^{x}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)\mathrm{Vol}(S_i) + \phi_\mu \sum_{i=x+1}^{h}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)\mathrm{Vol}(S_i) \\
&\overset{(c)}{=} \phi_0 \sum_{i=1}^{x}\mathbf{p}_i^2 d(v_i) - \phi_0 \mathbf{p}_{x+1}^2 \mathrm{Vol}(S_x) + \phi_\mu \mathbf{p}_{x+1}^2 \mathrm{Vol}(S_x) + \phi_\mu \sum_{i=x+1}^{h}\mathbf{p}_i^2 d(v_i) \\
&\overset{(d)}{=} \phi_\mu \sum_{i=1}^{h}\mathbf{p}_i^2 d(v_i) - (\phi_\mu - \phi_0)\sum_{i=1}^{x}(\mathbf{p}_i^2 - \mathbf{p}_{x+1}^2)d(v_i) \\
&\overset{(e)}{=} \phi_\mu \sum_{i=1}^{n}\mathbf{p}_i^2 d(v_i) - (\phi_\mu - \phi_0)\sum_{i=1}^{x+1}(\mathbf{p}_i^2 - \mathbf{p}_{x+1}^2)d(v_i)
\end{aligned} \tag{13}$$

where (a) is because $\mathbf{p}_{h+1} = \ldots = \mathbf{p}_n = 0$, (b) uses the definitions of ϕ_0, ϕ_μ and the fact (6), (c) and (d) rearrange terms, and lastly (e) again uses $\mathbf{p}_{h+1} = \ldots = \mathbf{p}_n = 0$.

We notice that $y+1$ has a symmetric role with $x+1$ because $x+1$ is the smallest integer such that $\sum_{i=1}^{x+1} d(v_i) = \mathrm{Vol}(S_{x+1}) \geq \mu \mathrm{Vol}(G)$, and $y+1$ is the largest integer that $\sum_{i=y+1}^{n} d(v_i) = \mathrm{Vol}(\bar{S}_y) \geq \mu \mathrm{Vol}(G)$. Together with the fact that $\mathbf{p}$ and $\mathbf{q}$ are symmetric, we get the following result

$$\sum_{i=1}^{n-1}\left(\mathbf{q}_{i+1}^2 - \mathbf{q}_i^2\right)|\partial S_i| \geq \phi_\mu \sum_{i=1}^{n}\mathbf{q}_i^2 d(v_i) - (\phi_\mu - \phi_0)\sum_{i=y+1}^{n}(\mathbf{q}_i^2 - \mathbf{q}_{y+1}^2)d(v_i). \tag{14}$$

Combining (13) and (14), we get that

$$\begin{aligned}
&\frac{\sum_{i=1}^{n-1}\left(\mathbf{p}_i^2 - \mathbf{p}_{i+1}^2\right)|\partial S_i| + \sum_{i=1}^{n-1}\left(\mathbf{q}_{i+1}^2 - \mathbf{q}_i^2\right)|\partial S_i|}{\sum_{v \in V} d(v)\left(\mathbf{p}_v^2 + \mathbf{q}_v^2\right)} \\
&\geq \phi_\mu - \left(\phi_\mu - \phi_0\right)\frac{\sum_{i=1}^{x+1}(\mathbf{p}_i^2 - \mathbf{p}_{x+1}^2)d(v_i) + \sum_{i=y+1}^{n}(\mathbf{q}_i^2 - \mathbf{q}_{y+1}^2)d(v_i)}{\sum_{i=1}^{n} d(v_i)\left(\mathbf{p}_i^2 + \mathbf{q}_i^2\right)} \\
&\overset{(f)}{\geq} \phi_\mu - \left(\phi_\mu - \phi_0\right)\frac{\sum_{i=1}^{x+1}(\mathbf{p}_i^2 - \mathbf{p}_{x+1}^2)d(v_i) + \sum_{i=y+1}^{n}(\mathbf{q}_i^2 - \mathbf{q}_{y+1}^2)d(v_i)}{\sum_{i=1}^{x+1}\mathbf{p}_i^2 d(v_i) + \sum_{i=y+2}^{n}\mathbf{q}_i^2 d(v_i)} \\
&= \phi_\mu - \left(\phi_\mu - \phi_0\right)\frac{\sum_{i=1}^{x+1}(\mathbf{p}_i^2 - \mathbf{p}_{x+1}^2)d(v_i) + \sum_{i=y+2}^{n}(\mathbf{q}_i^2 - \mathbf{q}_{y+1}^2)d(v_i)}{\sum_{i=1}^{x+1}\mathbf{p}_i^2 d(v_i) + \sum_{i=y+2}^{n}\mathbf{q}_i^2 d(v_i)} \\
&\overset{(g)}{=} \phi_0 + \left(\phi_\mu - \phi_0\right)\frac{\mathbf{p}_{x+1}^2 \mathrm{Vol}(S_{x+1}) + \mathbf{q}_{y+1}^2 \mathrm{Vol}(\bar{S}_{y+1})}{\sum_{i=1}^{x+1}\mathbf{p}_i^2 d(v_i) + \sum_{i=y+2}^{n}\mathbf{q}_i^2 d(v_i)},
\end{aligned} \tag{15}$$

where (f) is because $\sum_{i=x+2}^{n} \mathbf{p}_i^2 d(v_i) + \sum_{i=1}^{y+1} \mathbf{q}_i^2 d(v_i) \geq 0$ and (g) rearranges terms.

As shown in Lemma 3.3,

$$\mathbf{g}(v_{x+1}) \geq \sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}.$$

We also know that

$$-\sqrt{\frac{1-\mu}{\mu\mathrm{Vol}(G)}} \leq \mathbf{g}(v_i) \leq \sqrt{\frac{(1-\mu)}{\mu\mathrm{Vol}(G)}},$$

and

$$\mathbf{g}(v_{h+1}) \leq -\sqrt{\frac{\mu}{(1-\mu)\mathrm{Vol}(G)}}$$

because we made the assumption that $\mathbf{g}(v_{h+1}) < 0$. Therefore for any $i \leq x$,

$$\begin{aligned}\frac{\mathbf{p}_i}{\mathbf{p}_{x+1}} &= \frac{\mathbf{g}(v_i) - \mathbf{g}(v_{h+1})}{\mathbf{g}(v_{x+1}) - \mathbf{g}(v_{h+1})} = 1 + \frac{\mathbf{g}(v_i) - \mathbf{g}(v_{x+1})}{\mathbf{g}(v_{x+1}) - \mathbf{g}(v_{h+1})} \\ &\leq 1 + \left(\sqrt{\frac{1-\mu}{\mu}} - \sqrt{\frac{\mu}{1-\mu}}\right) / \left(2\sqrt{\frac{\mu}{1-\mu}}\right) = \frac{1}{2\mu}\end{aligned}$$

and for any $i \geq y+2$,

$$\frac{\mathbf{q}_i}{\mathbf{p}_{x+1}} = \frac{\mathbf{g}(v_{h+1}) - \mathbf{g}(v_i)}{\mathbf{g}(v_{x+1}) - \mathbf{g}(v_{h+1})} \leq \left(-\sqrt{\frac{\mu}{1-\mu}} + \sqrt{\frac{1-\mu}{\mu}}\right) / \left(2\sqrt{\frac{\mu}{1-\mu}}\right) = \frac{1}{2\mu} - 1.$$

Hence

$$\begin{aligned}&\frac{\mathbf{p}_{x+1}^2 \mathrm{Vol}(S_{x+1}) + \mathbf{q}_{y+2}^2 \mathrm{Vol}(\bar{S}_{y+1})}{\sum_{i=1}^{x+1} \mathbf{p}_i^2 d(v_i) + \sum_{i=y+2}^{n} \mathbf{q}_i^2 d(v_i)} \\ &\geq \frac{\mathbf{p}_{x+1}^2 \mathrm{Vol}(S_{x+1})}{\sum_{i=1}^{x+1} \mathbf{p}_i^2 d(v_i) + \sum_{i=y+2}^{n} \mathbf{q}_i^2 d(v_i)} \\ &= \frac{\mathrm{Vol}(S_{x+1})}{\sum_{i=1}^{x+1} d(v_i)\mathbf{p}_i^2/\mathbf{p}_{x+1}^2 + \sum_{i=y+2}^{n} d(v_i)\mathbf{q}_i^2/\mathbf{p}_{x+1}^2} \\ &\geq \frac{\mathrm{Vol}(S_{x+1})}{\mathrm{Vol}(S_{x+1})/4\mu^2 + (1/2\mu - 1)^2 \mathrm{Vol}(\bar{S}_{y+2})} \\ &= \frac{1}{1/4\mu^2 + (1/2\mu - 1)^2 \mathrm{Vol}(\bar{S}_{y+2})/\mathrm{Vol}(S_{x+1})} \\ &\overset{(h)}{\geq} \frac{1}{1/4\mu^2 + (1/2\mu - 1)^2} \\ &= \frac{\mu^2}{\mu^2 - \mu + 1/2},\end{aligned}$$

where (h) uses the fact that $\mathrm{Vol}(\bar{S}_{y+2}) \leq \mu\mathrm{Vol}(G) \leq \mathrm{Vol}(S_{x+1})$. Plugging the above inequality into (15) establishes (12), and hence the lower bound of Theorem 3.1.

4. Computational considerations and examples

Because of the nonconvex constraints, such as $\mathbf{x}^T \boldsymbol{D}\mathbf{x} = 1$, the program in (5) is nonconvex. In addition, the constraints on the magnitudes of the entries make the problem no longer a generalized eigenvalue problem. We are not aware of any technique that can compute the global optimum of this program.

One important use of this spectral program, as shown in Huang et al. [19], is to lower bound the μ-conductance; in other words, it corresponds to the easier direction of Theorem 3.1, $\phi_\mu \geq \lambda_\mu/2$. To track this lower bound effectively, we perform a sequence of transformations as originally shown in [19].

We first relax (5) to the following semidefinite program:

$$\begin{aligned} \lambda_\mu^{\mathrm{SDP}} = \underset{\boldsymbol{X}\in\mathbb{R}^{|V|\times|V|}}{\text{minimize}} \quad & \operatorname{Tr}(\boldsymbol{L}\boldsymbol{X}) \\ \text{subject to} \quad & \operatorname{Tr}(\boldsymbol{D}\boldsymbol{X}) = 1 \\ & \operatorname{Tr}\left(\mathbf{d}\mathbf{d}^T\boldsymbol{X}\right) = 0 \\ & \operatorname{Diag}(\boldsymbol{X}) \le \frac{(1-\mu)\mathbf{1}}{\mu\mathrm{Vol}(G)} \\ & \operatorname{Diag}(\boldsymbol{X}) \ge \frac{\mu\mathbf{1}}{(1-\mu)\mathrm{Vol}(G)}. \end{aligned} \tag{16}$$

This relaxation is obtained by lifting $\mathbf{x}\mathbf{x}^T$ to a positive semidefinite matrix $\boldsymbol{X}$, and it makes the program convex. As a consequence, $\lambda_\mu^{\mathrm{SDP}}/2$ is a valid lower bound on ϕ_μ.

Despite being convex, semidefinite programs do not scale well and are impractical for graphs with millions of vertices and tens of millions of edges. However, the above SDP is weakly constrained, with only $2|V|+2 = 2n+2$ constraints. By the result of Barvinok and Pataki [5,29], such programs admit an optimal solution of rank $O(\sqrt{C})$, where C is the number of constraints. In our case, this implies the existence of a solution of rank $O(\sqrt{n})$.

Motivated by this observation, we consider the classical Burer-Monteiro factorization [8–10], which factorizes the $n\times n$ decision variable $\boldsymbol{X}$ into the product $\boldsymbol{Y}\boldsymbol{Y}^T$ where $\boldsymbol{Y}\in\mathbb{R}^{n\times r}$. The factorization automatically eliminates the positive semidefinite constraint. The resulting low-rank semidefinite program can be written as

$$\begin{aligned} \lambda_\mu^{\mathrm{LRSDP}} = \underset{\boldsymbol{Y}\in\mathbb{R}^{n\times r}}{\text{minimize}} \quad & \operatorname{Tr}\left(\boldsymbol{L}\boldsymbol{Y}\boldsymbol{Y}^T\right) \\ \text{subject to} \quad & \operatorname{Tr}\left(\boldsymbol{D}\boldsymbol{Y}\boldsymbol{Y}^T\right) = 1 \\ & \operatorname{Tr}\left(\mathbf{d}\mathbf{d}^T\boldsymbol{Y}\boldsymbol{Y}^T\right) = 0 \\ & \operatorname{Diag}(\boldsymbol{Y}\boldsymbol{Y}^T) \le \frac{(1-\mu)\mathbf{1}}{\mu\mathrm{Vol}(G)} \\ & \operatorname{Diag}(\boldsymbol{Y}\boldsymbol{Y}^T) \ge \frac{\mu\mathbf{1}}{(1-\mu)\mathrm{Vol}(G)}. \end{aligned} \tag{17}$$

Although this factorization substantially reduces the number of variables, it renders the optimization problem nonconvex. Recent results show that, for generic problem instances, all local optima of low-rank semidefinite programs of this form are global optima as long as $r = \Omega(\sqrt{C})$ [7,14]. In practice, to further improve scalability, one may start with a rank $r \ll \sqrt{n}$ and gradually increase it when the primal–dual gap fails to improve [18]. Empirically, a small rank such as $r\approx 10$ is often sufficient to reach the global optimum on many problem instances. Even when the solution $\boldsymbol{Y}$ is not optimal, the dual objective still provides a valid lower bound on λ_μ.

For more computational examples along this line, where the spectral program and its semidefinite relaxation are used to certify lower bounds on μ-conductance on large real-world graphs, we refer the reader to Huang et al. [19].

Computational example for Theorem 3.1 On the other hand, here we provide one simple computational example to understand how the bound of Theorem 3.1 behaves on graphs in practice [2] Because tracking exact μ-conductance is NP-hard, we pick a small graph with only 85 vertices such that we can exactly solve the integer program of μ-conductance using `Gurobi`. To make the graph easier to visualize, we first generate the graph in a geometric way. We randomly sample vertices in two-dimensional space and add edges from each vertex to its nearest neighbors. This graph contains a cut with conductance ≈ 0.01 for a small set of vertices S. Then to simulate the phenomenon of many real-world graphs that μ-conductance gets much larger when μ gets closer to $1/2$, we add random edges into the graph in an Erdős-Rényi way. In particular, we add edges between each pair of vertices outside S with probability 0.2, and edges incident to S with probability 0.001 to preserve the low-conductance cut. This graph ends up having 568 edges, with an edge density 0.159. To compute the lower and upper bounds, we then compute ϕ_μ, ϕ_0 exactly using the integer program (4). To track λ_μ, we solve the semidefinite program relaxation (16) using the package `CSDP`. The results are summarized in Fig. 1.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: David F. Gleich reports financial support was provided by National Science Foundation. David F. Gleich

[2] Code for this example is available at https://github.com/luotuoqingshan/mu_cond_cheeger.

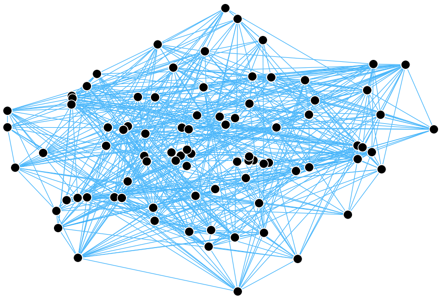

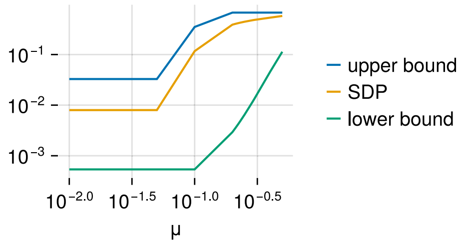

Fig. 1. An example graph with 85 vertices and 568 edges (Left) and the bounds from Lemma 3.2 computed on it for different μs (Right). To keep track of λ_μ, the optimum of the spectral program, we compute its relaxation $\lambda_\mu^{\mathrm{SDP}}$. The true λ_μ lies between the blue upper bound and the orange SDP relaxation. We see that the lower bound gets tighter in the large μ regime. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

reports financial support was provided by US Department of Energy. David F. Gleich reports financial support was provided by Intelligence Advanced Research Projects Activity. David F. Gleich reports a relationship with FlexCompute that includes: consulting or advisory. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data and code is made public through a github repo.

References

[1] Noga Alon, Vitali D. Milman, λ_1, isoperimetric inequalities for graphs, and superconcentrators, J. Combin. Theory Ser. B 38 (1) (1985) 73–88.
[2] Konstantin Andreev, Harald Racke, Balanced graph partitioning, Theory Comput. Syst. 39 (6) (Oct. 2006) 929–939.
[3] Sanjeev Arora, Satish Rao, Umesh Vazirani, Expander flows, geometric embeddings and graph partitioning, J. ACM 56 (2) (Apr. 2009) 1–37.
[4] Nikhil Bansal, et al., Min-max graph partitioning and small set expansion, SIAM J. Comput. 43 (2) (Jan. 2014) 872–904.
[5] A.I. Barvinok, Problems of distance geometry and convex properties of quadratic maps, Discrete Comput. Geom. 13 (2) (Mar. 1995) 189–202.
[6] Austin R. Benson, David F. Gleich, Jure Leskovec, Higher-order organization of complex networks, Science 353 (6295) (July 2016) 163–166.
[7] Nicolas Boumal, Vladislav Voroninski, Afonso S. Bandeira, The non-convex Burer-Monteiro approach works on smooth semidefinite programs, in: Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, 2016, pp. 2757–2765.
[8] Samuel Burer, Changhui Choi, Computational enhancements in low-rank semidefinite programming, Optim. Methods Softw. 21 (3) (June 2006) 493–512.
[9] Samuel Burer, Renato D.C. Monteiro, A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization, Math. Program. 95 (2) (Feb. 2003) 329–357.
[10] Samuel Burer, Renato D.C. Monteiro, Local minima and convergence in low-rank semidefinite programming, Math. Program. 103 (3) (Dec. 2004) 427–444.
[11] Jeff Cheeger, A lower bound for the smallest eigenvalue of the Laplacian, in: Proceedings of the Princeton Conference in Honor of Professor S. Bochner, 1969, pp. 195–199, English (US).
[12] F. Chung, Four proofs of the Cheeger inequality and graph partition algorithms, in: Proceedings of the ICCM, 2007, pp. 751–772.
[13] Fan Chung, Spectral Graph Theory, American Mathematical Society, Dec. 1996.
[14] Diego Cifuentes, On the Burer–Monteiro method for general semidefinite programs, Optim. Lett. 15 (6) (Jan. 2021) 2299–2309.
[15] Jozef Dodziuk, Difference equations, isoperimetric inequality and transience of certain random walks, Trans. Amer. Math. Soc. 284 (2) (1984) 787–794.
[16] Lester Randolph Ford, D.R. Fulkerson, Flows in Networks, Princeton University Press, Dec. 1963.
[17] Rumi Ghosh, et al., The interplay between dynamics and networks: centrality, communities, and Cheeger inequality, in: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, ACM, New York, New York, USA, 2014, pp. 1406–1415.
[18] Yufan Huang, David F. Gleich, Suboptimality bounds for trace-bounded sdps enable a faster and scalable low-rank sdp solver sdplr+, arXiv preprint, arXiv:2406.10407, 2024.
[19] Yufan Huang, C. Seshadhri, David F. Gleich, Theoretical bounds on the network community profile from low-rank semi-definite programming, in: Andreas Krause, et al. (Eds.), Proceedings of the 40th International Conference on Machine Learning, in: Proceedings of Machine Learning Research, vol. 202, PMLR, 23–29 Jul 2023, pp. 13976–13992.
[20] Rohit Khandekar, Satish Rao, Umesh Vazirani, Graph partitioning using single commodity flows, J. ACM 56 (4) (June 2009) 1–15.
[21] Ioannis Koutis, Gary Miller, Richard Peng, A generalized Cheeger inequality, Linear Algebra Appl. 665 (May 2023) 139–152.
[22] Tsz Chiu Kwok, et al., Improved Cheeger's inequality: analysis of spectral partitioning algorithms through higher order spectral gap, in: Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC'13, ACM, June 2013, pp. 11–20.
[23] Lap Chi Lau, Kam Chuen Tung, Robert Wang, Cheeger inequalities for directed graphs and hypergraphs using reweighted eigenvalues, in: Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC '23, ACM, June 2023, pp. 1834–1847.
[24] James R. Lee, Shayan Oveis Gharan, Luca Trevisan, Multiway spectral partitioning and higher-order Cheeger inequalities, J. ACM 61 (6) (Dec. 2014) 1–30.
[25] Tom Leighton, Satish Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, J. ACM 46 (6) (Nov. 1999) 787–832.
[26] Jure Leskovec, et al., Community structure in large networks: natural cluster sizes and the absence of large well-defined clusters, Internet Math. 6 (1) (Jan. 2009) 29–123.
[27] Anand Louis, et al., Many sparse cuts via higher eigenvalues, in: Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, STOC'12, ACM, May 2012, pp. 1131–1140.

[28] L. Lovász, M. Simonovits, The mixing rate of Markov chains, an isoperimetric inequality, and computing the volume, in: Proceedings of the 31st Annual Symposium on Foundations of Computer Science, vol. 1, SFCS '90, IEEE Computer Society, Washington, DC, USA, 1990, pp. 346–354.
[29] Gábor Pataki, On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues, Math. Oper. Res. 23 (2) (May 1998) 339–358.
[30] Satu Elisa Schaeffer, Graph clustering, Comput. Sci. Rev. 1 (1) (Aug. 2007) 27–64.
[31] Jiří Šíma, Satu Elisa Schaeffer, On the NP-completeness of some graph cluster measures, in: SOFSEM 2006: Theory and Practice of Computer Science, in: Lecture Notes in Computer Science, vol. 3831, 2006, pp. 530–537.
[32] Yao Zhu, David F. Gleich, A parallel min-cut algorithm using iteratively reweighted least squares, Parallel Comput. 59 (Nov. 2016) 43–59.