

Graph Reduction with Spectral and Cut Guarantees

Andreas Loukas

ANDREAS.LOUKAS@EPFL.CH

*Laboratoire de traitement des signaux 2,
École Polytechnique Fédérale de Lausanne,
CH-1015 Lausanne, Switzerland*

Editor: Karsten Borgwardt

Abstract

Can one reduce the size of a graph without significantly altering its basic properties? The graph reduction problem is hereby approached from the perspective of *restricted spectral approximation*, a modification of the spectral similarity measure used for graph sparsification. This choice is motivated by the observation that restricted approximation carries strong *spectral* and *cut* guarantees, and that it implies approximation results for unsupervised learning problems relying on spectral embeddings. The article then focuses on coarsening—the most common type of graph reduction. Sufficient conditions are derived for a small graph to approximate a larger one in the sense of restricted approximation. These findings give rise to algorithms that, compared to both standard and advanced graph reduction methods, find coarse graphs of improved quality, often by a large margin, without sacrificing speed.

Keywords: graph reduction and coarsening, spectral methods, unsupervised learning

1. Introduction

As graphs grow in size, it becomes pertinent to look for generic ways of simplifying their structure while preserving key properties. Simplified graph representations find profound use in the design of approximation algorithms, can facilitate storage and retrieval, and ultimately ease graph data analysis by separating overall trends from details.

There are two main ways to simplify graphs. First, one may reduce the number of edges—a technique commonly referred to as *graph sparsification*. In a series of works, it has been shown that it is possible to find sparse graphs that approximate all pairwise distances (Peg and Schäffer, 1989), every cut (Karger, 1999), or every eigenvalue (Spielman and Teng, 2011)—respectively referred to as spanners, cut sparsifiers and spectral sparsifiers. Spectral sparsification techniques, in particular, can yield computational benefits whenever the number of edges is the main bottleneck (Batson et al., 2013). Indeed, they form a fundamental component of nearly-linear time algorithms for linear systems involving symmetric diagonally dominant matrices (Koutis et al., 2010; Spielman and Srivastava, 2011), and have found application to machine learning problems involving graph-structured data (Candriello et al., 2018).

Alternatively, one may seek to reduce the size of the graph directly, i.e., the number of its vertices N , by some form of vertex selection or re-combination scheme followed by re-wiring. This idea can be traced back to the multigrid literature, that targets the acceleration of finite-element methods using cycles of multi-level coarsening, lifting, and refinement. After being generalized to graphs, reduction methods have become pervasive in computer science and form a key element of modern graph processing pipelines, especially with regards to graph partitioning (Hendrickson and Leland, 1995; Karypis and Kumar, 1998; Kushnir et al., 2006; Dhillon et al., 2007; Wang et al., 2014) and graph visualization (Koren, 2002; Hu, 2005; Walshaw, 2006). In machine learning, reduction methods are used to create multi-scale representations of graph-structured data (Lafon and Lee, 2006; Gavish et al., 2010; Shuman et al., 2016) and as a layer of graph convolutional neural networks (Bruna et al., 2014; Defferrard et al., 2016; Bronstein et al., 2017; Simonovsky and Komodakis, 2017; Ardizzone et al., 2018; Liang et al., 2018). In addition, being shown to solve linear systems in (empirically) linear time (Koutis et al., 2011; Livne and Brandt, 2012) as well as to approximate the Fiedler vector (Urschel et al., 2014; Gandhi, 2016), reduction methods have been used to accelerate inverse problems employing graph-regularization (Hirani et al., 2015; Colley et al., 2017). Some of their main benefits are the ability to deal with sparse graphs (graphs with at most $O(N \log N)$ edges) and to accelerate algorithms whose complexity depends on the number of vertices as well as edges.

In contrast to graph sparsification, there has been only circumstantial theory supporting graph reduction (Moitra, 2011; Dörfler and Bullo, 2013; Loukas and Vandergheynst, 2018). The lack of a concrete understanding of how different reduction choices affect fundamental graph properties is an issue: the significant majority of reduction algorithms in modern graph processing and machine learning pipelines have been designed based on intuition and possess no rigorous justification or provable guarantees.

1.1. A Spectral Perspective

My starting point in this work is *spectral similarity*—a measure that has been proven useful in sparsification for determining how well a graph approximates another one. To render spectral similarity applicable to graphs of different sizes, I generalize it and restrict it over a subspace of size that is at most equal to the size of the reduced graph. I refer to the resulting definition as *restricted spectral approximation*¹ (or restricted approximation for short). Despite being a statement about subspaces, restricted approximation has significant consequences. It is shown that when the subspace in question is a principal eigenspace (this is a feature agnostic choice where one wants to preserve only the graph structure), the eigenvalues and eigenspaces of the reduced graph approximate those of the original large graph. It is then a corollary that (i) if the large graph has a good (normalized) cut so does the smaller one; and (ii) that unsupervised learning algorithms that utilize spectral embeddings, such as spectral clustering (Von Luxburg, 2007) and Laplacian eigenmaps (Belkin and Niyogi, 2003), can also work well when ran on the smaller graph and their solution is lifted. More

1. Though similarly named, the definition of *restricted spectral similarity* previously proposed by Loukas and Vandergheynst (2018) concerns a set of vectors (rather than subspaces) and is significantly weaker than the one examined here.

generally, it is expected that methods relying on specific eigenspaces will exhibit similar behavior when combined with coarsening, e.g., in graph signal regression (Loukas and Perraudin, 2016; Grassi et al., 2017), compression (Shahid et al., 2016), sampling (Puy et al., 2018), or forecasting (Isufi et al., 2018).

The analysis then focuses on *graph coarsening*—a popular type of reduction where, in each level, reduced vertices are formed by contracting disjoint sets of connected vertices (each such set is called a *contraction set*). I derive sufficient conditions for a small coarse graph to approximate a larger graph in the sense of restricted spectral approximation. Crucially, this result holds for any number of levels and is independent of how the subspace is chosen. Though the derived bound is global, a decoupling argument renders it locally separable over levels and contraction sets, facilitating computation. The final bound can be interpreted as measuring the *local variation* over each contraction set, as it involves the maximum variation of vectors supported on each contracted subgraph.

These findings give rise to algorithms for graph coarsening that I refer to as *local variation algorithms*. Each such algorithm starts from a predefined family of candidate contraction sets. Even though any connected set of vertices may form a valid candidate set, I opt for small well-connected sets constructed based on simple rules: in the *edge-based variant* one candidate set is constructed for each of pair of adjacent vertices, whereas in the *neighborhood-based variant* every candidate set contains a vertex along with all adjacent vertices. The algorithm then greedily² contracts those sets whose local variation is the smallest. Depending on how the candidate family is constructed, the proposed algorithms obtain different solutions, trading off computational complexity for reduction.

1.2. Theoretical and Practical Implications

This work improves and generalizes upon the state-of-the-art in several ways:

- Instead of directly focusing on specific constructions, a general graph reduction scheme is studied featuring coarsening as a special case. As a consequence, the implications of restricted approximation are proven in a fairly general setting where specifics of the reduction, such as the type of graph representation and the reduction matrices involved, are abstracted.
- Contrary to previous results on the analysis of coarsening (Loukas and Vandergheynst, 2018), the analysis holds for multiple levels of reduction. Given that the majority of coarsening methods reduce the number of vertices by a constant factor at each level, a multi-level approach is necessary to achieve significant reduction. Along that line, the analysis also brings an intuitive insight: rather than taking the standard approach of approximating at each level the graph produced by the previous level, one should strive to preserve the properties of the original graph at every level.

2. Even after decoupling, the problem of candidate set selection is not only NP-hard but also cannot be approximated to a constant factor in polynomial time (by reduction to the maximum-weight independent set problem). For the specific case of edge-based families, where one candidate set is constructed for each pair of adjacent vertices, the greedy iterative contraction can be substituted by more sophisticated procedures accompanied by improved guarantees.

- The proposed local variation algorithms are not heuristically designed, but optimize (an upper bound of) the restricted spectral approximation objective. Despite the breadth of the literature that utilizes some form of graph reduction and coarsening, the overwhelming majority of known methods are heuristics (Safro et al., 2015). A notable exception is Kron reduction (Dörfler and Bullo, 2013), an elegant method that aims to preserve the effective resistance distance. Compared to Kron reduction, the methods proposed here are accompanied by stronger spectral guarantees (i.e., beyond interlacing), do not sacrifice the sparsity of the graph, and can ultimately be more scalable as they do not rely on the Schur complement of the Laplacian matrix.

To demonstrate the practical benefits of local variation methods, the analysis is complemented with numerical results on representative graphs ranging from scale-free to planar graphs. Compared to both standard (Karypis and Kumar, 1998) and advanced reduction methods (Ron et al., 2011; Livne and Brandt, 2012; Shuman et al., 2016), the proposed methods yield small graphs of improved spectral quality, often by a large margin, without being much slower than naive heavy-edge matching. A case in point: when examining how close are the principal eigenvalues of the coarse and original graph for a reduction of 70%, local variation methods attain on average $2.6\times$ smaller error; this gain becomes $3.9\times$ if one does not include Kron reduction in the comparison.

2. Graph Reduction and Coarsening

The following section introduces graph reduction. The exposition starts by considering a general reduction scheme. It is then shown how graph coarsening arises if one additionally imposes few natural restrictions on the interpretability of reduced variables.

2.1. Graph Reduction

Consider a positive semidefinite (PSD) matrix $L \in \mathbb{R}^{N \times N}$ whose sparsity structure captures the connectivity structure of a connected weighted undirected graph $G = (\mathcal{V}, \mathcal{E}, W)$ of $N = |\mathcal{V}|$ vertices and $M = |\mathcal{E}|$ edges. In other words, $L(i, j) \neq 0$ only if e_{ij} is a valid edge between vertices v_i and v_j . Moreover, let x be an arbitrary vector of size N .

I study the following generic reduction scheme:

Scheme 1: Graph reduction

Commence by setting $L_0 = L$ and $x_0 = x$ and proceed according to the following two recursive equations:

$$L_\ell = P_\ell^\mp L_{\ell-1} P_\ell^+ \quad \text{and} \quad x_\ell = P_\ell x_{\ell-1},$$

where $P_\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ are matrices with more columns than rows, $\ell = 1, 2, \dots, c$ is the level of the reduction, symbol $+$ (resp. \mp) denotes the pseudoinverse (resp. transposed

Scheme 1: Graph reduction (cont.)

pseudoinverse), and N_ℓ is the dimensionality at level ℓ such that $N_0 = N$ and $N_c = n \ll N$.

Vector x_c is lifted back to \mathbb{R}^N by recursion $\tilde{x}_{\ell-1} = P_\ell^+ \tilde{x}_\ell$, where $\tilde{x}_c = x_c$.

Graph reduction thus involves a sequence of $c + 1$ graphs

$$G = G_0 = (\mathcal{V}_0, \mathcal{E}_0, W_0) \quad G_1 = (\mathcal{V}_1, \mathcal{E}_1, W_1) \quad \cdots \quad G_c = (\mathcal{V}_c, \mathcal{E}_c, W_c) \quad (1)$$

of decreasing size $N = N_0 > N_1 > \cdots > N_c = n$, where the sparsity structure of L_ℓ matches that of graph G_ℓ , and each vertex of G_ℓ represents one of more vertices of $G_{\ell-1}$.

The multi-level design allows us to achieve high *dimensionality reduction ratio*

$$r = 1 - \frac{n}{N},$$

even when at each level the ratio $r_\ell = 1 - \frac{N_\ell}{N_{\ell-1}}$ is small. For instance, supposing that $r_\ell \geq \varrho$ for each ℓ , then $c = O(\log(n/N)/\log(1 - \varrho))$ levels suffice to reduce the dimension to n .

One may express the reduced quantities in a more compact form:

$$x_c = Px, \quad L_c = P^\top LP^+ \quad \text{and} \quad \tilde{x} = \Pi x, \quad (2)$$

where $P = P_c \cdots P_1$, $P^+ = P_1^+ \cdots P_c^+$, $\Pi = P^+ P$, and for convenience, I drop zero indices and refer to a lifted vector as $\tilde{x} (= \tilde{x}_0)$. The rationale of this scheme is that vector \tilde{x} should be the best approximation of x given P in an ℓ_2 -sense, which is a consequence of the following property:

Property 1 Π is a projection matrix.

On the other hand, matrix L is reduced such that $x_c^\top L_c x_c = \tilde{x}^\top L \tilde{x}$.

Though introduced here for the reduction of sparse PSD matrices representing the sparsity structure of a graph, Scheme 1 can also be applied to any PSD matrix L . This and similar reduction schemes belong to the class of Nyström methods and, to the extent of my knowledge, were first studied in the context of approximate low-rank matrix approximation (Halko et al., 2011; Wang and Zhang, 2013). Despite the common starting point, interpreting L and L_c as sparse “graph matrices”, as it is done here, incorporates a graph-theoretic twist to reduction, distinguishing it from previous approaches³: the constructions that we will study are eventually more scalable and interpretable as they maintain the graph structure of L after reduction. Obtaining guarantees is also arguably more challenging in this setting, as the involved problems end up being strongly combinatorial.

3. To achieve low-rank approximation, matrix P is usually built by sampling columns of L .

2.2. Properties of Reduced Graphs

Even in the general context where P is an arbitrary $n \times N$ matrix, certain handy properties can be proven about the relation between L_c and L .

To begin with, it is simple to see that the set of positive semidefinite matrices is closed under reduction.

Property 2 *If L is PSD, then so is L_c .*

The proof is elementary: if L is PSD then there exists matrix S such that $L = S^\top S$, implying that $L_c = P^\top L P^+$ can also be written as $L_c = S_c^\top S_c$ if one sets $S_c = S P^+$.

I further consider the spectrum of the two matrices. Sort the eigenvalues of L as $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$ and denote by $\tilde{\lambda}_k$ the k -th largest eigenvalue of L_c and \tilde{u}_k the associated eigenvector.

It turns out that the eigenvalues $\tilde{\lambda}$ and λ are interlaced.

Theorem 3 *For any P with full-row rank and $k = 1, \dots, n$, we have*

$$\gamma_1 \lambda_k \leq \tilde{\lambda}_k \leq \gamma_2 \lambda_{k+N-n}$$

with $\gamma_1 = \lambda_1((PP^\top)^{-1})$ and $\gamma_2 = \lambda_n((PP^\top)^{-1})$, respectively the smallest and largest eigenvalue of $(PP^\top)^{-1}$.

The above result is a generalization of the Cauchy interlacing theorem for the case that $PP^\top \neq I$. It also resembles the interlacing inequalities known for the normalized Laplacian, where the re-normalization is obtained by construction. Chen et al. (2004) showed in Theorem 2.7 of their paper that after contracting $N - n$ edges $\lambda_{k-N+n} \leq \lambda_k \leq \lambda_{k+N-n}$ for $k = 1, 2, \dots, n$ and with $\lambda_\ell = 0$ when $\ell \leq 0$, resembling the upper bound above. The lower bound is akin to that given by Chung (1997, Lemma 1.15), again for the normalized Laplacian. Also notably, the inequalities are similar to those known for Kron reduction (Dörfler and Bullo, 2013, Lemma 3.6).

Theorem 3 is particularly pessimistic as it has to hold for every possible P and L (subject to γ_1 and γ_2). Much stronger results will be obtained later on by restricting the attention to constructions that satisfy additional properties (see Theorem 13).

One can also say something about the action of L_c on vectors.

Property 4 *For every vector $x \in \text{im}(\Pi)$, one has*

$$x_c^\top L_c x_c = x^\top \Pi L \Pi x = x^\top L x \quad \text{and} \quad \tilde{x} = \Pi x = x.$$

In other words, reduction maintains the action of L of every vector that lies in the image of Π . Naturally, after lifting the eigenvectors \tilde{u}_k of L_c are included in this set.

2.3. Coarsening as a Type of Graph Reduction

Coarsening is a type of graph reduction abiding to a set of constraints that render the graph transformation interpretable. More precisely, in coarsening one selects for each level ℓ a surjective (i.e., many-to-one) map $\varphi_\ell : \mathcal{V}_{\ell-1} \rightarrow \mathcal{V}_\ell$ between the original vertex set $\mathcal{V}_{\ell-1}$ and the smaller vertex set \mathcal{V}_ℓ . I refer to the set of vertices $\mathcal{V}_{\ell-1}^{(r)} \subseteq \mathcal{V}_{\ell-1}$ mapped onto the same vertex v'_r of \mathcal{V}_ℓ as a *contraction set*:

$$\mathcal{V}_{\ell-1}^{(r)} = \{v \in \mathcal{V}_{\ell-1} : \varphi_\ell(v) = v'_r\}$$

For a graphical depiction of contraction sets, see Figure 1. I also constrain φ_ℓ slightly by requiring that the subgraph of $G_{\ell-1}$ induced by each contraction set $\mathcal{V}_{\ell-1}^{(r)}$ is connected.

It is easy to deduce that contraction sets induce a partitioning of $\mathcal{V}_{\ell-1}$ into N_ℓ subgraphs, each corresponding to a single vertex of \mathcal{V}_ℓ . Every reduced variable thus corresponds to a small set of adjacent vertices in the original graph and coarsening amounts to a scaling operation. An appropriately constructed coarse graph aims to capture the global problem structure, whereas neglected details can be recovered in a local refinement phase.

Coarsening can be placed in the context of Scheme 1 by restricting each P_ℓ to lie in the family of coarsening matrices, defined next:

Definition 5 (Coarsening matrix) *Matrix $P_\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ is a coarsening matrix w.r.t. graph $G_{\ell-1}$ if and only if it satisfies the following two conditions:*

- a. *It is a surjective mapping of the vertex set, meaning that if $P_\ell(r, i) \neq 0$ then $P_\ell(r', i) = 0$ for every $r' \neq r$.*
- b. *It is locality preserving, equivalently, the subgraph of $G_{\ell-1}$ induced by the non-zero entries of $P_\ell(r, :)$ is connected for each r .*

An interesting consequence of this definition is that, in contrast to graph reduction, with coarsening matrices the expensive pseudo-inverse computation can be substituted by simple transposition and re-scaling:

Proposition 6 (Easy inversion) *The pseudo-inverse of a coarsening matrix P_ℓ is given by $P_\ell^+ = P_\ell^\top D_\ell^{-2}$, where D_ℓ is the diagonal matrix with $D_\ell(r, r) = \|P_\ell(r, :)\|_2$.*

Proposition 6 carries two consequences. First, coarsening can be done in linear time. Each coarsening level (both in the forward and backward directions) entails multiplication by a sparse matrix. Furthermore, both P_ℓ and P_ℓ^+ have only $N_{\ell-1}$ non-zero entries meaning that $O(N)$ and $O(M)$ operations suffice to coarsen respectively a vector and a matrix L whose sparsity structure reflects the graph connectivity. In addition, the number of graph edges also decreases at each level. Denoting by μ_ℓ the average number of edges of the graphs induced by contraction sets $\mathcal{V}_{\ell-1}^{(r)}$ for every r , then a quick calculation reveals that the coarsest graph has $m = M - \sum_{\ell=1}^c N_\ell \mu_\ell$ edges. If, for instance, at each level all vertices are perfectly contracted into pairs then $\mu_\ell = 2$ and $N_\ell = N/2^\ell$, meaning that $m = M - 2N(1 - 2^{-c})$.




Figure 1: Toy coarsening example. Grey discs denote contraction sets. The first three vertices of G forming contraction set \mathcal{V}_0^1 are contracted onto vertex v'_1 . All other vertices remain unaffected.

2.4. Laplacian Consistent Coarsening

A further restriction that can be imposed is that coarsening is consistent w.r.t. the Laplacian form. Suppose that L is the combinatorial Laplacian of G defined as

$$L(i, j) = \begin{cases} d_i & \text{if } i = j \\ -w_{ij} & \text{if } e_{ij} \in \mathcal{E} \\ 0 & \text{otherwise,} \end{cases}$$

where w_{ij} is the weight associated with edge e_{ij} and d_i the weighted degree of v_i . The following proposition can be proven:

Proposition 7 (Consistency) *Let P be a coarsening matrix w.r.t. a graph with combinatorial Laplacian L . Matrix $L_c = P^\top L P^+$ is a combinatorial Laplacian if and only if the non-zero entries of P^+ are equally valued.*

It is a corollary of Propositions 6 and 7 that in Laplacian consistent coarsening, for any $v'_r \in \mathcal{V}_\ell$ and $v_i \in \mathcal{V}_{\ell-1}$ matrices $P_\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and $P_\ell^+ \in \mathbb{R}^{N_{\ell-1} \times N_\ell}$ are given by:

$$P_\ell(r, i) = \begin{cases} \frac{1}{|\mathcal{V}_{\ell-1}^{(r)}|} & \text{if } v_i \in \mathcal{V}_{\ell-1}^{(r)} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad [P_\ell^+](i, r) = \begin{cases} 1 & \text{if } v_i \in \mathcal{V}_{\ell-1}^{(r)} \\ 0 & \text{otherwise,} \end{cases}$$

where the contraction sets $\mathcal{V}_{\ell-1}^{(1)}, \dots, \mathcal{V}_{\ell-1}^{(N_\ell)}$ were defined in Section 2.3.

2.4.1. A TOY EXAMPLE

The toy graph shown in Figure 1a whose Laplacian is given by

$$L = \begin{bmatrix} 3 & -1 & -1 & -1 & 0 \\ -1 & 3 & -1 & 0 & -1 \\ -1 & -1 & 2 & 0 & 0 \\ -1 & 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 & 1 \end{bmatrix}$$

illustrates an example where the gray vertices $\mathcal{V}_0^{(1)} = \{v_1, v_2, v_3\}$ of G are contracted onto vertex v'_1 , as shown in Figure 1b. The main matrices I have defined are

$$P_1 = \begin{bmatrix} 1/3 & 1/3 & 1/3 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad P_1^+ = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \Pi = P_1^+ P_1 = \begin{bmatrix} 1/3 & 1/3 & 1/3 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

and coarsening results in

$$L_c = P_1^\top L P_1^+ = \begin{bmatrix} 2 & -1 & -1 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix} \quad x_c = P_1 x = \begin{bmatrix} (x(1) + x(2) + x(3))/3 \\ x(4) \\ x(5) \end{bmatrix}.$$

Finally, when lifted x_c becomes

$$\tilde{x} = P_1^+ x_c = \begin{bmatrix} (x(1) + x(2) + x(3))/3 \\ (x(1) + x(2) + x(3))/3 \\ (x(1) + x(2) + x(3))/3 \\ x(4) \\ x(5) \end{bmatrix}.$$

Since vertices v_4 and v_5 are not affected, the respective contraction sets $\mathcal{V}_0^{(2)}$ and $\mathcal{V}_0^{(3)}$ are singleton sets.

2.5. Properties of Laplacian Consistent Coarsening

Due to its particular construction, Laplacian consistent coarsening is accompanied some interesting properties. I present three in the following:

Cuts. To begin with, weights of edges in G_c correspond to weights of cuts in G .

Property 8 For any level ℓ , the weight $W_\ell(r, q)$ between vertices $v'_r, v'_q \in \mathcal{V}_\ell$ is equal to

$$W_\ell(r, q) = \sum_{v_i \in \mathcal{S}_\ell^{(r)}} \sum_{v_j \in \mathcal{S}_\ell^{(q)}} w_{ij},$$

where $\mathcal{S}_\ell^{(r)} = \{v_i \in \mathcal{V} : \varphi_\ell \circ \dots \circ \varphi_1(v_i) = v'_r\} \subset \mathcal{V}$ contains all vertices of G contracted onto $v'_r \in \mathcal{V}_\ell$.

In the toy example, there exists a single edge of unit weight connecting vertices in $\mathcal{V}_0^{(1)}$ and $\mathcal{V}_0^{(2)}$, and as such the weight between v'_1 and v'_2 is equal to one.

Eigenvalue interlacing. For a single level of Laplacian consistent coarsening, matrix $PP^\top = P_1 P_1^\top$ is given by $\text{diag}(1/|\mathcal{V}_0^{(1)}|, \dots, 1/|\mathcal{V}_0^{(N_1)}|)$, implying that the multiplicative constants in Theorem 3 are:

$$\gamma_1 = \min_{v_i \in \mathcal{V}} |\mathcal{V}_0^{\varphi_1(v_i)}| \geq 1 \quad \text{and} \quad \gamma_2 = \max_{v_i \in \mathcal{V}} |\mathcal{V}_0^{\varphi_1(v_i)}|.$$

Above, $v'_r = \varphi_1(v_i) \in \mathcal{V}_1$ is the vertex to which v_i is mapped to and the set $\mathcal{V}_0^{\varphi_1(v_i)}$ contains all vertices also contracted to v'_r . Thus in the toy example, $\lambda_k \leq \tilde{\lambda}_k \leq 3\lambda_{k+2}$ for every $k \leq 3$. If multiple levels are utilized these terms become dependent on the sequence of contractions. To obtain a general bound let $\varphi_1^\ell(v_i) = \varphi_\ell \circ \dots \circ \varphi_1(v_i) \in \mathcal{V}_\ell$ be the vertex onto which $v_i \in \mathcal{V}$ is contracted to in the ℓ -th level.

Property 9 *If L_c is obtained from L by Laplacian consistent coarsening, then*

$$\gamma_1 \geq \min_{v_i \in \mathcal{V}} \prod_{\ell=1}^c |\mathcal{V}_{\ell-1}^{\varphi_1^\ell(v_i)}| \geq 1 \quad \text{and} \quad \gamma_2 \leq \max_{v_i \in \mathcal{V}} \prod_{\ell=1}^c |\mathcal{V}_{\ell-1}^{\varphi_1^\ell(v_i)}|,$$

with the set $\mathcal{V}_{\ell-1}^{\varphi_1^\ell(v_i)}$ containing all vertices of $\mathcal{V}_{\ell-1}$ that are contracted onto $\varphi_1^\ell(v_i)$.

The proof follows from the diagonal form of $P_\ell \dots P_1 P_1^\top \dots P_\ell^\top$ and the special row structure of each P_ℓ for every ℓ , but it's not included for brevity. The dependency of $\tilde{\lambda}_k$ on the size of contraction sets can be removed either by enforcing at each level that all contraction sets have identical size and dividing the graph weights by that size, or by re-normalizing each P_ℓ such that $P_\ell^\top = P_\ell^+$. The latter approach was used by Loukas and Vandergheynst (2018) but is not adopted here as it results in L_c losing its Laplacian form.

Nullspace. Finally, as is desirable, the structure of the nullspace of L is preserved both by coarsening and lifting:

Property 10 *If P is a (multi-level) Laplacian consistent coarsening matrix, then*

$$P1_N = 1_n \quad \text{and} \quad P^+1_n = 1_N,$$

where the subscript indicates the dimensionality of the constant vector.

Thus, we can casually ignore vectors parallel to the constant vector in our analysis.

3. Restricted Notions of Approximation

This section formalizes how should a graph be reduced such that fundamental structural properties (e.g., its spectrum and cuts) are preserved. Inspired by work in graph sparsification, I introduce a measure of approximation that is tailored to graph reduction. The new definition implies strong guarantees about the distance of the original and coarsened spectrum and gives conditions such that the cut structure of a graph is preserved by coarsening.

3.1. Restricted Spectral Approximation

One way to define how close a PSD matrix L is to its reduced counterpart is to establish an isometry guarantee w.r.t. the following induced semi-norms:

$$\|x\|_L = \sqrt{x^\top L x} \quad \text{and} \quad \|x_c\|_{L_c} = \sqrt{x_c^\top L_c x_c}$$

Ideally, one would hope that there exists $\epsilon > 0$ such that

$$(1 - \epsilon) \|x\|_L \leq \|x_c\|_{L_c} \leq (1 + \epsilon) \|x\|_L \quad (3)$$

for all $x \in \mathbb{R}^N$.

If the inequalities in (3) hold, matrices L_c and L are called ϵ -similar. The objective of constructing sparse spectrally similar graphs is the main idea of spectral graph sparsifiers, a popular method for accelerating the solution of linear systems involving the Laplacian. In addition, spectral similarity carries several interesting consequences that are of great help in the construction of approximation algorithms: the eigenvalues and eigenvectors of two similar graphs are close and, moreover, all vertex partitions have similar cut size.

In contrast to graph sparsification however, since here the dimension of the space changes it is impossible to satisfy (3) for every $x \in \mathbb{R}^N$ unless one trivially sets $\epsilon = 1$ (this follows by a simple rank argument). To carry out a meaningful analysis, one needs to consider a subspace of dimension $k \leq n$ and aim to approximate the behavior of L solely within it.

With this in mind, I define the following generalization of spectral similarity:

Definition 11 (Restricted spectral approximation) *Let \mathbf{R} be a k -dimensional subspace of \mathbb{R}^N . Matrices L_c and L are (\mathbf{R}, ϵ) -similar if there exists an $\epsilon \geq 0$ such that*

$$\|x - \tilde{x}\|_L \leq \epsilon \|x\|_L \quad \text{for all } x \in \mathbf{R},$$

where $\tilde{x} = P^+Px$.

In addition to the restriction on \mathbf{R} , the above definition differs from (3) in how the error is measured. In fact, it asserts a property that is slightly stronger than an approximate isometry w.r.t. a semi-norm within \mathbf{R} . The strengthening of the notion of approximation deviates from the restricted spectral similarity property proposed by Loukas and Vandergheynst (2018) and is a key ingredient in obtaining multi-level bounds. Nevertheless, one may recover a restricted spectral similarity-type guarantee as a direct consequence:

Corollary 12 *If L_c and L are (\mathbf{R}, ϵ) -similar, then*

$$(1 - \epsilon) \|x\|_L \leq \|x_c\|_{L_c} \leq (1 + \epsilon) \|x\|_L \quad \text{for all } x \in \mathbf{R}.$$

Proof Let S be defined such that $L = S^\top S$. By the triangle inequality:

$$|\|x\|_L - \|x_c\|_{L_c}| = |\|Sx\|_2 - \|SP^+Px\|_2| \leq \|Sx - SP^+Px\|_2 = \|x - \tilde{x}\|_L \leq \epsilon \|x\|_L,$$

which is equivalent to the claimed relation. ■

Clearly, if L_c and L are (\mathbf{R}, ϵ) -similar then they are also (\mathbf{R}', ϵ') -similar, where \mathbf{R}' is any subspace of \mathbf{R} and $\epsilon' \geq \epsilon$. As such, results about large subspaces and small ϵ are the most desirable.

It will be shown in Sections 3.2 and 3.3 that the above definition implies restricted versions of the spectral and cut guarantees provided by spectral similarity. For instance, instead of attempting to approximate the entire spectrum as done by spectral graph sparsifiers, here one can focus on a subset of the spectrum with particular significance.

3.2. Implications for the Graph Spectrum

One of the primary benefits of restricted spectral approximation is that it implies a relation between the spectra of matrices L and L_c that goes beyond interlacing (see Theorem 3).

To this effect, consider the smallest k eigenvalues and corresponding eigenvectors and define the following matrices:

$$U_k \in \mathbb{R}^{N \times k} = [u_1, u_2, \dots, u_k] \quad \text{and} \quad \Lambda_k = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_k)$$

As I will show next, ensuring that ϵ in Proposition 17 is small when $\mathbf{R} = \mathbf{U}_k \triangleq \text{span}(U_k)$ suffices to guarantee that the first k eigenvalues and eigenvectors of L and L_c are aligned.

The first result concerns eigenvalues.

Theorem 13 (Eigenvalue approximation) *If L_c and L are $(\mathbf{U}_k, \epsilon_k)$ -similar, then*

$$\gamma_1 \lambda_k \leq \tilde{\lambda}_k \leq \gamma_2 \frac{(1 + \epsilon_k)^2}{1 - \epsilon_k^2 (\lambda_k / \lambda_2)} \lambda_k,$$

whenever $\epsilon_k^2 < \lambda_2 / \lambda_k$.

Crucially, the bound depends on λ_k instead of λ_{k+N-n} and thus can be significantly tighter than the one given by Theorem 3. Noticing that $\epsilon_k \leq \epsilon_{k'}$ whenever $k < k'$, one also deduces that it is stronger for smaller eigenvalues. For $k = 2$ in particular, one has

$$\gamma_1 \lambda_2 \leq \tilde{\lambda}_2 \leq \gamma_2 \frac{(1 + \epsilon_2)^2}{1 - \epsilon_2^2} \lambda_2,$$

which can be made small by controlling ϵ_2 .

I also analyze the angle between principal eigenspaces of L and L_c . I follow Li (1994) and split the eigendecompositions of $L = U \Lambda U^\top$ and $P^\top L_c P = P^\top \tilde{U} \tilde{\Lambda} \tilde{U}^\top P$ as

$$L = (U_k, U_{k^\perp}) \begin{pmatrix} \Lambda_k & \\ & \Lambda_{k^\perp} \end{pmatrix} \begin{pmatrix} U_k^\top \\ U_{k^\perp}^\top \end{pmatrix} \quad P^\top L_c P = (P^\top \tilde{U}_k, P^\top \tilde{U}_{k^\perp}) \begin{pmatrix} \tilde{\Lambda}_k & \\ & \tilde{\Lambda}_{k^\perp} \end{pmatrix} \begin{pmatrix} \tilde{U}_k^\top P \\ \tilde{U}_{k^\perp}^\top P \end{pmatrix},$$

where $\tilde{\Lambda}_k$ and \tilde{U}_k are defined analogously to Λ_k and U_k . Davis and Kahan (1970) defined the *canonical angles* between the spaces spanned by U_k and $P^\top \tilde{U}_k$ as the singular values of the matrix

$$\Theta(U_k, P^\top \tilde{U}_k) \triangleq \arccos(U_k^\top P^\top \tilde{U}_k \tilde{U}_k^\top P U_k)^{-1/2},$$

see also (Stewart, 1990). The smaller the sinus of the canonical angles are the closer the two subspaces lie. The following theorem reveals a connection between the Frobenius norm of the sinus of the canonical angles and restricted spectral approximation.

Theorem 14 (Eigenspace approximation) *If L_c and L are $(\mathbf{U}_k, \epsilon_k)$ -similar then*

$$\left\| \sin \Theta(U_k, P^\top \tilde{U}_k) \right\|_F^2 \leq \frac{1}{\lambda_{k+1} - \lambda_k} \left(\sum_{i \leq k} \lambda_i \left(\frac{(1 + \epsilon_i)^2}{\gamma_1} - 1 \right) + \lambda_k \sum_{i \leq k} \epsilon_i \right),$$

Note that the theorem above utilizes all ϵ_i with $i \leq k$, corresponding to the restricted spectral approximation constants for $\mathbf{R} = \mathbf{U}_i$, respectively. However, all these can be trivially relaxed to ϵ_k , since $\epsilon_i \leq \epsilon_k$ for all $i \leq k$.

3.3. Implications for Graph Partitioning

One of the most popular applications of coarsening is to accelerate graph partitioning (Hendrickson and Leland, 1995; Karypis and Kumar, 1998; Kushnir et al., 2006; Dhillon et al., 2007; Wang et al., 2014). In the following, I provide a rigorous justification for this choice by showing that if the (Laplacian consistent) coarsening is done well and G_c contains a good normalized cut, then so will G . For the specific case of spectral clustering, I also provide an explicit bound on the coarse solution quality.

3.3.1. EXISTENCE RESULTS

For consistent coarsening, the spectrum approximation results presented previously imply similarities between the cut-structures of G_c and G .

To formalize this intuition, the conductance of any subset \mathcal{S} of \mathcal{V} is defined as

$$\phi(\mathcal{S}) \triangleq \frac{w(\mathcal{S}, \bar{\mathcal{S}})}{\min\{w(\mathcal{S}), w(\bar{\mathcal{S}})\}},$$

where $\bar{\mathcal{S}} = \mathcal{V} \setminus \mathcal{S}$ is the complement set, $w(\mathcal{S}, \bar{\mathcal{S}}) = \sum_{v_i \in \mathcal{S}, v_j \in \bar{\mathcal{S}}} w_{ij}$ is the weight of the cut and $w(\mathcal{S}) = \sum_{v_i \in \mathcal{S}} \sum_{v_j \in \mathcal{V}} w_{ij}$ is the volume of \mathcal{S} .

The k -conductance of a graph measures how easy it is to cut it into k disjoint subsets $\mathcal{S}_1, \dots, \mathcal{S}_k \subset \mathcal{V}$ of balanced volume:

$$\phi_k(G) = \min_{\mathcal{S}_1, \dots, \mathcal{S}_k} \max_i \phi(\mathcal{S}_i),$$

where the minimum is taken over all possible k -partitionings $(\mathcal{S}_1, \dots, \mathcal{S}_k)$. The smaller $\phi_k(G)$ is, the easier it is to partition our graph.

As it turns out, restricted spectral approximation can be used to relate the conductance of the original and coarse graphs. To state the result, it will be useful to denote by D the diagonal degree matrix and further to suppose that V_k contains the first k eigenvectors of the normalized Laplacian $L^n = D^{-1/2} L D^{-1/2}$, whose eigenvalues are $0 = \mu_1 \leq \dots \leq \mu_k$.

Theorem 15 For any graph G and integer $2 \leq k \leq \lfloor n/2 \rfloor$, if L_c and L are $(\mathbf{R}_{2k}, \epsilon_{2k})$ -similar combinatorial Laplacian matrices then

$$\phi_k(G) \leq \phi_k(G_c) = O\left(\sqrt{\frac{\gamma_2(1 + \epsilon_{2k})^2 \xi_k(G)}{1 - \epsilon_{2k}^2(\mu_{2k}/\mu_2)}} \phi_k(G)\right)$$

with $\mathbf{R}_{2k} = \text{span}(D^{-1/2}V_{2k})$ and $\xi_k(G) \leq \log k$, whenever $\epsilon_{2k}^2 < \mu_2/\mu_{2k}$. If G is planar then $\xi_k(G) \leq 1$ and if G excludes K_h as a minor (but is not planar) then $\xi_k(G) \leq h^4$. For $k = 2$, supposing that L_c and L are $(\mathbf{R}_2, \epsilon_2)$ -similar, we additionally have

$$\phi_2(G) \leq \phi_2(G_c) \leq 2\sqrt{\frac{\gamma_2(1 + \epsilon_2)^2}{1 - \epsilon_2^2}} \phi_2(G).$$

The above theorem gives a non-constructive result: it does not reveal how to find the optimal partitioning, but provides conditions such that the latter is of similar quality in the two graphs.

3.3.2. SPECTRAL CLUSTERING

It is also possible to derive approximation results about the solution quality of unsupervised learning algorithms that utilize the first k eigenvectors in order to partition G . I focus here on spectral clustering. To perform the analysis, let U_k and $P^\top \tilde{U}_k$ be the spectral embedding of the vertices w.r.t. L and L_c , respectively, and define the optimal partitioning as

$$\mathcal{P}^* = \arg \min_{\mathcal{P}=\{\mathcal{S}_1, \dots, \mathcal{S}_k\}} \mathcal{F}_k(U_k, \mathcal{P}) \quad \text{and} \quad \tilde{\mathcal{P}}^* = \arg \min_{\mathcal{P}=\{\mathcal{S}_1, \dots, \mathcal{S}_k\}} \mathcal{F}_k(P^\top \tilde{U}_k, \mathcal{P}), \quad (4)$$

where, for any embedding X , the k -means cost induced by partitioning \mathcal{V} into clusters $\mathcal{S}_1, \dots, \mathcal{S}_k$ is defined as

$$\mathcal{F}_k(X, \mathcal{P}) \triangleq \sum_{z=1}^k \sum_{v_i, v_j \in \mathcal{S}_z} \frac{\|X(i, \cdot) - X(j, \cdot)\|_2^2}{2|\mathcal{S}_z|}.$$

One then measures the quality of $\tilde{\mathcal{P}}^*$ by examining how far the correct minimizer $\mathcal{F}_k(U_k, \mathcal{P}^*)$ is to $\mathcal{F}_k(U_k, \tilde{\mathcal{P}}^*)$. Boutsidis et al. (2015) noted that if the two quantities are close then, despite the clusters themselves possibly being different, they both feature the same quality w.r.t. the k -means objective.

An end-to-end control of the k -means error is obtained by combining the inequality derived by Loukas and Vandergheynst (2018), based on the works of (Boutsidis et al., 2015; Yu et al., 2014; Martin et al., 2018), $|\mathcal{F}_k(U_k, \mathcal{P}^*)^{1/2} - \mathcal{F}_k(U_k, \tilde{\mathcal{P}}^*)^{1/2}| \leq 2\sqrt{2} \left\| \sin \Theta(U_k, P^\top \tilde{U}_k) \right\|_F$ with Theorem 14:

Corollary 16 *If L_c and L are $(\mathbf{U}_k, \epsilon_k)$ -similar then*

$$\left(\mathcal{F}_k(U_k, \mathcal{P}^*)^{1/2} - \mathcal{F}_k(U_k, \tilde{\mathcal{P}}^*)^{1/2}\right)^2 \leq \frac{8}{\lambda_{k+1} - \lambda_k} \left(\sum_{i \leq k} \lambda_i \left(\frac{(1 + \epsilon_i)^2}{\gamma_1} - 1 \right) + \lambda_k \sum_{i \leq k} \epsilon_i \right).$$

Contrary to previous analysis (Loukas and Vandergheynst, 2018), the approximation result here is applicable to any number of levels and it can be adapted to hold for the eigenvectors of the normalized Laplacian⁴. Nevertheless, it should be stressed that at this point it is an open question whether the above analysis yields benefits over other approaches tailored especially to the acceleration of spectral clustering. A plethora of such specialized algorithms are known (Tremblay and Loukas, 2019)—arguing about the pros and cons of each extends beyond the scope of this work.

4. Graph Coarsening by Local Variation

This section proposes algorithms for Laplacian consistent coarsening. I suppose that L is a combinatorial graph Laplacian and, given a subspace \mathbf{R} and a target graph size n , aim to find an (\mathbf{R}, ϵ) -similar Laplacian L_c of size $n \times n$ with ϵ smaller than some threshold ϵ' .

Local variation algorithms differ only in the type of contraction sets that they consider. For instance, the edge-based local variation algorithm only contracts edges, whereas in the neighborhood-based variant each contraction set is a subset of the neighborhood of a vertex. Otherwise, all local variation algorithms follow the same general methodology and aim to minimize an upper bound of ϵ . To this end, two bounds are exploited: First, L_c is shown to be (\mathbf{R}, ϵ) -similar to L with $\epsilon \leq \prod_{\ell} (1 + \sigma_{\ell}) - 1$, where the *variation cost* σ_{ℓ} depends only on previous levels (see Section 4.1). The main difficulty with minimizing σ_{ℓ} is that it depends on interactions between contraction sets. For this reason, the second bound shows that these interactions can be decoupled by considering each *local variation cost*, i.e., the cost of contracting solely the vertices in $\mathcal{V}_{\ell-1}^{(r)}$, independently on a slightly modified subgraph (see Section 4.2). Having achieved this, Section 4.3 considers ways of efficiently identifying disjoint contraction sets with small local variation cost.

4.1. Decoupling Levels and the Variation Cost

Guaranteeing restricted spectral approximation w.r.t. subspace \mathbf{R} boils down to minimizing at each level ℓ the *variation cost*

$$\sigma_{\ell} = \|\Pi_{\ell}^{\perp} A_{\ell-1}\|_{L_{\ell-1}} = \|S_{\ell-1} \Pi_{\ell}^{\perp} A_{\ell-1}\|_2,$$

where $L_{\ell-1} = S_{\ell-1}^{\top} S_{\ell-1}$ and $\Pi_{\ell}^{\perp} = I - P_{\ell}^{\dagger} P_{\ell}$ is a projection matrix. Matrix $A_{\ell-1}$ captures two types of information:

4. For the normalized Laplacian, one should perform (combinatorial) Laplacian consistent coarsening on a modified eigenspace, as in the proof of Theorem 15.

1. Foremost, it encodes the behavior of the target Laplacian matrix L w.r.t. \mathbf{R} . This is clearly seen in the first level, for which one has that $A_0 = VV^\top L^{+1/2}$ with $V \in \mathbb{R}^{N \times k}$ being an orthonormal basis of \mathbf{R} .
2. When $\ell > 1$ one needs to consider A_0 in view of the reduction done in previous levels. The necessary modification turns out to be $A_{\ell-1} = B_{\ell-1}(B_{\ell-1}^\top L_{\ell-1} B_{\ell-1})^{+1/2}$, with $B_{\ell-1} = P_{\ell-1} B_{\ell-2} \in \mathbb{R}^{N_{\ell-1} \times N}$ expressed in a recursive manner and $B_0 = A_0$.

The following result makes explicit the connection between ϵ and σ_ℓ .

Proposition 17 *Matrices L_c and L are (\mathbf{R}, ϵ) -similar with $\epsilon \leq \prod_{\ell=1}^c (1 + \sigma_\ell) - 1$.*

Crucially, the previous result makes it possible to design a multi-level coarsening greedily, by starting from the first level and optimizing following levels one at a time. Thus, every local variation algorithm operates in the following manner:

Algorithm 1 Multi-level coarsening

- 1: **input:** Combinatorial Laplacian L , threshold ϵ' , and target size n .
 - 2: Set $\ell \leftarrow 0$, $L_\ell \leftarrow L$, and $\epsilon_\ell \leftarrow 0$.
 - 3: **while** $N_\ell > n$ and $\epsilon_\ell < \epsilon'$ **do**
 - 4: $\ell \leftarrow \ell + 1$
 - 5: Coarsen $L_{\ell-1}$ using Algorithm 2 with threshold $\sigma' = \frac{1+\epsilon'}{1+\epsilon_{\ell-1}} - 1$ and target size n .
 Let L_ℓ be the resulting Laplacian of size N_ℓ with variation cost σ_ℓ .
 - 6: $\epsilon_\ell \leftarrow (1 + \epsilon_{\ell-1})(1 + \sigma_\ell) - 1$.
 - 7: **return** L_ℓ
-

Algorithm 1 returns a Laplacian matrix L_c that is (\mathbf{R}, ϵ) -similar to L with $\epsilon \leq \epsilon_c \leq \epsilon'$, where c is the last level ℓ . On the other hand, setting ϵ' to a large value ensures that the same algorithm always attains the target reduction at the expense of loose restricted approximation guarantees.

Remark. The variation cost simplifies when \mathbf{R} is an eigenspace of L . I demonstrate this for the choice of \mathbf{U}_k , though an identical argument can be easily derived for any eigenspace. Denote by Λ the diagonal $N \times N$ eigenvalue matrix placed from top-left to bottom-right in non-decreasing order and by U the respective full eigenvector matrix. Furthermore, let Λ_k be the $k \times k$ sub-matrix of Λ with the smallest k eigenvalues in its diagonal. By the unitary invariance of the spectral norm, it follows that $\sigma_0 = \|\Pi_1^\perp U_k U_k^\top L^{+1/2}\|_{L_0} = \|\Pi_1^\perp U_k U_k^\top L^{+1/2} U\|_{L_0} = \|\Pi_1^\perp U_k U_k^\top U \Lambda^{+1/2}\|_{L_0}$. Simplifying and eliminating zero columns, one may redefine $B_0 = U_k \Lambda_k^{+1/2} \in \mathbb{R}^{N \times k}$, such that once more $\sigma_0 = \|\Pi_1^\perp B_0\|_{L_0}$. This is computationally attractive because now at each level one needs to take the pseudo-inverse-square-root of a $k \times k$ matrix $B_{\ell-1}^\top L_{\ell-1} B_{\ell-1}$, with $k \ll N$.

4.2. Decoupling Contraction Sets and Local Variation

Suppose that $\Pi_{\mathcal{C}}^\perp$ is the (complement) projection matrix obtained by contracting solely the vertices in set \mathcal{C} , while leaving all other vertices in $\mathcal{V}_{\ell-1}$ untouched:

$$\left[\Pi_{\mathcal{C}}^\perp x \right] (i) = \begin{cases} x(i) - \sum_{v_j \in \mathcal{C}} \frac{x(j)}{|\mathcal{C}|} & \text{if } v_i \in \mathcal{C} \\ 0 & \text{otherwise.} \end{cases}$$

(Here, for convenience, the level index is suppressed.)

Furthermore, let $L_{\mathcal{C}}$ be the $N_{\ell-1} \times N_{\ell-1}$ combinatorial Laplacian whose weight matrix is

$$[W_{\mathcal{C}}] (i, j) = \begin{cases} W_{\ell-1}(i, j) & \text{if } v_i, v_j \in \mathcal{C} \\ 2 W_{\ell-1}(i, j) & \text{if } v_i \in \mathcal{C} \text{ and } v_j \notin \mathcal{C} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

That is, $W_{\mathcal{C}}$ is zero everywhere other than at the edges touching at least one vertex in \mathcal{C} . The following proposition shows us how to decouple the contribution of each contraction set to the variation cost.

Proposition 18 *The variation cost is bounded by*

$$\sigma_\ell^2 \leq \sum_{\mathcal{C} \in \mathcal{P}_\ell} \|\Pi_{\mathcal{C}}^\perp A_{\ell-1}\|_{L_{\mathcal{C}}}^2,$$

where $\mathcal{P}_\ell = \{\mathcal{V}_{\ell-1}^{(1)}, \dots, \mathcal{V}_{\ell-1}^{(N_\ell)}\}$ is the family of contraction sets of level ℓ .

The above argument therefore entails bounding the, difficult to optimize, variation cost as a function of locally computable and independent costs $\|\Pi_{\mathcal{C}}^\perp A_{\ell-1}\|_{L_{\mathcal{C}}}^2$. The obtained expression is a relaxation, as it assumes that the interaction between contraction sets will be the worst possible. It might be interesting to notice that the quality of the relaxation depends on the weight of the cut between contraction sets. Taking the limit, the inequality converges to an equality as the weight of the cut shrinks. Also of note, the bound becomes tighter the larger the per-level dimensionality reduction requested (the smaller $N_\ell = |\mathcal{P}_\ell|$ is, the fewer inequalities are involved in the derivation).

4.3. Local Variation Coarsening Algorithms

Starting from a *candidate family* $\mathcal{F}_\ell = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots\}$, that is, an appropriately sized family of candidate contraction sets, the strategy will be to search for a small *contraction family* $\mathcal{P}_\ell = \{\mathcal{V}_{\ell-1}^{(1)}, \dots, \mathcal{V}_{\ell-1}^{(N_\ell)}\}$ with minimal variation cost σ_ℓ (\mathcal{P}_ℓ is valid if it partitions $\mathcal{V}_{\ell-1}$ into N_ℓ contraction sets). Every coarse vertex $v_r' \in \mathcal{V}_\ell$ is then formed by contracting the vertices in $\mathcal{V}_{\ell-1}^{(r)}$.

As a thought experiment, suppose that set $\mathcal{C} \in \mathcal{F}_\ell$ is chosen to be part of \mathcal{P}_ℓ . From the decoupling argument, its contribution to σ_ℓ^2 will be at most $\|\Pi_{\mathcal{C}}^\perp A_{\ell-1}\|_{L_{\mathcal{C}}}^2$ independently of

how other candidate sets are chosen. Moreover, the selection will reduce $N_{\ell-1}$ by $|\mathcal{C}| - 1$ vertices. Thus, one needs to look for the non-singleton candidate sets \mathcal{C} with cost

$$\text{cost}_\ell(\mathcal{C}) \triangleq \frac{\|\Pi_{\mathcal{C}}^\perp A_{\ell-1}\|_{L_{\mathcal{C}}}^2}{|\mathcal{C}| - 1} \quad (6)$$

that is as small as possible. I refer to (6) as *local variation cost* because it captures the maximal variation of all signals from an appropriate subspace (implied by $A_{\ell-1}$) with support on \mathcal{C} . On the other hand, since any permissible contraction family \mathcal{P}_ℓ should be a partitioning of $\mathcal{V}_{\ell-1}$, choosing \mathcal{C} precludes us from selecting any \mathcal{C}' with which it intersects.

Based on this intuition, Algorithm 2 sequentially examines candidate sets from \mathcal{F}_ℓ , starting from those with minimal cost. To decide whether a candidate set \mathcal{C} will be added to \mathcal{P}_ℓ the algorithm asserts that all vertices in \mathcal{C} are unmarked—essentially enforcing that all contraction sets are disjoint. Accordingly, as soon as \mathcal{C} is added to \mathcal{P}_ℓ , all vertices that are in \mathcal{C} become marked. Candidate sets with marked vertices are pruned ($\mathcal{C}' \leftarrow \mathcal{C} \setminus \text{marked}$) and their cost is updated. The algorithm terminates if either the target reduction is achieved, the error threshold is exceeded, or no candidate sets remain. Even though this remains implicit in the discussion, if at termination \mathcal{P}_ℓ does not cover every vertex of $\mathcal{V}_{\ell-1}$, then I compliment it with singleton sets, featuring one vertex each (and zero cost).

Algorithm 2 Single-level coarsening by local variation

- 1: **input:** Combinatorial Laplacian $L_{\ell-1}$, threshold σ' , and target size n .
 - 2: Form the family of candidate sets $\mathcal{F}_\ell = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots\}$ (algorithm-specific step).
 - 3: $N_\ell \leftarrow N_{\ell-1}$, **marked** $\leftarrow \emptyset$, $\sigma_\ell^2 \leftarrow 0$.
 - 4: Sort \mathcal{F}_ℓ in terms of increasing $\text{cost}_\ell(\mathcal{C})$.
 - 5: **while** $|\mathcal{F}_\ell| > 0$ and $N_\ell > n$ and $\sigma_\ell \leq \sigma'$ **do**
 - 6: Pop the candidate set \mathcal{C} of minimal cost s from \mathcal{F}_ℓ .
 - 7: **if** all vertices of \mathcal{C} are not **marked** and $\sigma' \geq \sqrt{\sigma_\ell^2 + (|\mathcal{C}| - 1)s}$ **then**
 - 8: **marked** $\leftarrow \text{marked} \cup \mathcal{C}$, $\mathcal{P}_\ell \leftarrow \mathcal{P}_\ell \cup \mathcal{C}$, $N_\ell \leftarrow N_\ell - |\mathcal{C}| + 1$, $\sigma_\ell^2 \leftarrow \sigma_\ell^2 + (|\mathcal{C}| - 1)s$
 - 9: **else**
 - 10: $\mathcal{C}' \leftarrow \mathcal{C} \setminus \text{marked}$
 - 11: **if** $|\mathcal{C}'| > 1$ **then**
 - 12: Compute $\text{cost}_\ell(\mathcal{C}')$ and insert \mathcal{C}' into \mathcal{F}_ℓ while keeping the latter sorted.
 - 13: Form the $N_\ell \times N_{\ell-1}$ coarsening matrix P_ℓ based on \mathcal{P}_ℓ .
 - 14: **return** $L_\ell \leftarrow P_\ell^\mp L_{\ell-1} P_\ell^+$ and σ_ℓ
-

Undeniably, Algorithm 2 is only one of the possible ways to select a partitioning of small variation cost. However, this algorithm stands out from other algorithms I experimented with, as it is efficient when the subspace of interest is an eigenspace (e.g., $V = U_k$), k is small, and the families \mathcal{F}_ℓ have been selected appropriately. Denote by $\Phi = \max_{\mathcal{C} \in \mathcal{F}_\ell} |\mathcal{C}|$ the maximum number of vertices in all candidate sets and by $\delta = \max_{\mathcal{C} \in \mathcal{F}_\ell} |\mathcal{C}|$ the cardinality of the maximum candidate set—I refer to these measures as *family weight* and *width*, respectively. Choosing $\mathbf{R} = \mathbf{U}_k$, the computational complexity of Algorithm 2 is $\tilde{O}(ckM + k^2N + ck^3 + c\Phi (\min\{k^2\delta + k\delta^2, k\delta^2 + \delta^3\} + \log \max_{\mathcal{C} \in \mathcal{F}_\ell} |\mathcal{C}|))$, which up to poly-log factors is linear on the number of edges, vertices, and Φ (see Appendix B for details).

If computational complexity is of no concern, one may consider the following two more sophisticated algorithms: *The optimal algorithm.* Given a candidate family, the algorithm that optimally minimizes the sum of local variation costs constructs a graph with one vertex for each subset of a candidate set and adds an edge between every two vertices whose respective sets have a non-empty intersection. It then selects \mathcal{P}_ℓ as the maximum independent set of minimal weight (the weight of each vertex is a local variation cost w.r.t a set). Unfortunately, even if the size of this graph is polynomial in N , this problem cannot be solved efficiently, since the minimum-weight independent set problem is NP-hard. Nevertheless, for the specific case where candidate sets correspond to edges the problem simplifies to a minimum-weight matching problem, which can be computed in $O(N_{\ell-1}^3)$ time exactly, whereas a $(2 + \delta)$ -approximation can be found much faster (Paz and Schwartzman, 2017). *The quadratic variant.* A second possibility is to proceed as with Algorithm 2, but to prune each $\mathcal{C}' \in \mathcal{F}_\ell$ after a set \mathcal{C} is added to \mathcal{P}_ℓ . The numerical experiments indicated that this additional step may improve the coarsening quality slightly, but it is not recommended for large graphs as it introduces a quadratic dependency of the complexity on N .

4.3.1. CANDIDATE CONTRACTION FAMILIES

To keep coarsening efficient, I focus on families of linear weight and almost constant width. Two possibilities are considered:

Edge-based. Here \mathcal{F}_ℓ contains one candidate set for each edge of $G_{\ell-1}$. This is a natural choice for coarsening—indeed, most coarsening algorithms in the literature use some form of edge contraction. It is straightforward to see that in this case $\Phi = 2M$ and $\delta = 2$, meaning that the expression of the computational complexity simplifies to $\tilde{O}(ckM + ck^3 + k^2N)$. The drawback of contracting edges is that at each level the graph size can only be reduced by at most a factor of 2, meaning that a large number of levels is necessary to achieve significant reduction⁵.

Neighborhood-based. A more attractive choice is to construct one candidate set for the neighborhood of each vertex, including the vertex itself. Denoting by Δ the largest combinatorial degree and since $\Phi = 2M$, the complexity here is bounded by $\tilde{O}(cM(k + \min\{k^2\Delta + k\Delta^2, k\Delta^2 + \Delta^3\}) + ck^3 + k^2N)$. The experiments show that the neighborhood-based variant generally achieves better reduction while being marginally slower than the edge-based variant.

As a final remark, when G is dense, the dependency on M can be dropped by sparsifying the graph before using Algorithm 1.

5. Numerical Results

The evaluation was performed on four representative graphs, each exhibiting different structural characteristics:

5. In practice, depending on the graph in question, the per-level reduction ratio r_ℓ is usually between 0.35 and 0.45.

- *Yeast*. Protein-to-protein interaction network in budding yeast, analyzed by Jeong et al. (2001). The network has $N = 1458$ vertices, $M = 1948$ edges, diameter of 19, and degree between 1 and 56.
- *Airfoil*. Finite-element graph obtained by airflow simulation Preis and Diekmann (1997), consisting of $N = 4000$ vertices, $M = 11490$ edges, diameter of 65, and degree between 1 and 9.
- *Minnesota*. Road network with $N = 2642$ vertices, $M = 3304$ edges, diameter of 99, and degree between 1 and 5 (Gleich, 2008).
- *Bunny*. Point cloud consisting of $N = 2503$ vertices, $M = 65490$ edges, diameter of 15, and degree between 13 and 97 (Turk and Levoy, 1994). The point cloud has been sub-sampled from its original size.

I compare to the following methods for multi-level graph reduction:

- *Heavy edge matching*. At each level of the scheme, the contraction family is obtained by computing a maximum-weight matching with the weight of each contraction set (v_i, v_j) calculated as $w_{ij} / \max\{d_i, d_j\}$. In this manner, heavier edges connecting vertices that are well separated from the rest of the graph are contracted first. Heavy edge matching was first introduced in the algebraic multigrid literature and, perhaps due to its simplicity, its variants have been repeatedly used for partitioning (Karypis and Kumar, 1998; Dhillon et al., 2007) and drawing (Walshaw, 2000; Hu, 2005) graphs, as well as more recently in graph convolutional neural networks (Defferrard et al., 2016).
- *Algebraic distance*. This method differs from heavy edge matching in that the weight of each candidate set $(v_i, v_j) \in \mathcal{E}$ is calculated as $(\sum_{q=1}^Q (x_q(i) - x_q(j))^2)^{1/2}$, where x_k is an N -dimensional test vector computed by successive sweeps of Jacobi relaxation. The complete method is described by Ron et al. (2011), see also (Chen and Safro, 2011). As recommended by the authors, I performed 20 relaxation sweeps. Further, I set the number of test vectors Q to equal the dimension k of the space I aimed to approximate (a simple rank argument shows that $Q \geq k$ for the test vectors to span the space).
- *Affinity*. This is a vertex proximity heuristic in the spirit of the algebraic distance that was proposed by Livne and Brandt (2012) in the context of their work on the lean algebraic multigrid. As per the author suggestion, the $Q = k$ test vectors are here computed by a single sweep of a Gauss-Seidel iteration.
- *Kron reduction*. At each level, the graph size is reduced by selecting a set of vertices of size $N/2$ (corresponding to the positive entries of the last eigenvector of L) and applying Kron reduction. The method, which was proposed by Shuman et al. (2016), is not strictly a coarsening method as it rewires the vertices of the reduced graph, resulting in significantly denser graphs⁶. Unfortunately, the rewiring step en-

6. As suggested by the authors, the sparsity of reduced graphs can be controlled by spectral sparsification. The sparsification step was not included in the numerical experiments since it often resulted in increased errors.

tails finding the Schur complement of a large Laplacian submatrix and thus generally exhibits $O(N^3)$ complexity, rendering it prohibitive for graphs of more than a few thousand vertices. Despite these drawbacks, the method is quite popular because of its elegant theoretical guarantees (Dörfler and Bullo, 2013).

Depending on how the edge matching is constructed, different variants of edge contraction methods can be implemented. At the two extremes of the complexity spectrum one finds the maximum matching of minimum weight at $O(N^3)$ complexity (Galil, 1986) or greedily constructs a matching by visiting vertices in random order and inducing $O(M)$ overhead (Dhillon et al., 2007).

For consistency, I implemented all edge-based methods by combining Algorithm 2 with an edge-based family and substituting the local variation cost with the (negative) method-specific edge weight. This generally yields matchings of better quality (heavier weight) than visiting vertices in a predefined order, at the expense of the marginally larger $O(M \log M)$ complexity necessary for sorting the edge weights. The choice is also motivated by the observation that the computational bottleneck of (sophisticated) edge contraction methods lies in the edge weight computation. For all experiments, I set $\epsilon' = \infty$ aiming for a fixed reduction rather than a restricted spectral approximation guarantee. The code reproducing the experiments can be accessed online ⁷.

5.1. Restricted Spectral Approximation

The first experiment tests how well L_c approximates the action of L w.r.t. the subspace U_k of the smallest variation. In other words, for each method, I plot the smallest ϵ such that the following equation holds:

$$\|x - \tilde{x}\|_L \leq \epsilon \|x\|_L \quad \text{for all } x \in \mathbf{U}_k. \tag{7}$$

The results are summarized in Figure 2 for two representative subspaces of size $k = 10$ and $k = 40$. With the exception of the Kron reduction that was repeated 10 times, all methods are deterministic and thus were run only once.

Overall, it can be seen that local variation methods outperform other coarsening methods in many problem instances. The gap is particularly prominent for large reductions, where multiple levels are employed. Neighborhood-based contraction yields the best result overall, mainly because it achieves the same reduction in fewer levels. Interestingly, local variation (and coarsening) methods in many cases also outperform Kron reduction, even though the latter is more demanding computationally.

I elaborate further on four points stemming from the results:

In most instances, a reduction of up to 70% is feasible while maintaining a decent approximation. The attained approximation is a function of the graph in question and k . Nevertheless, in almost all experiments, the best coarsening method could reduce the graph size by at least 70% while also ensuring that $\epsilon < 1$ (horizontal black line). This is an

⁷. github.com/loukasa/graph-coarsening/tree/v1.1 (DOI 175851068)

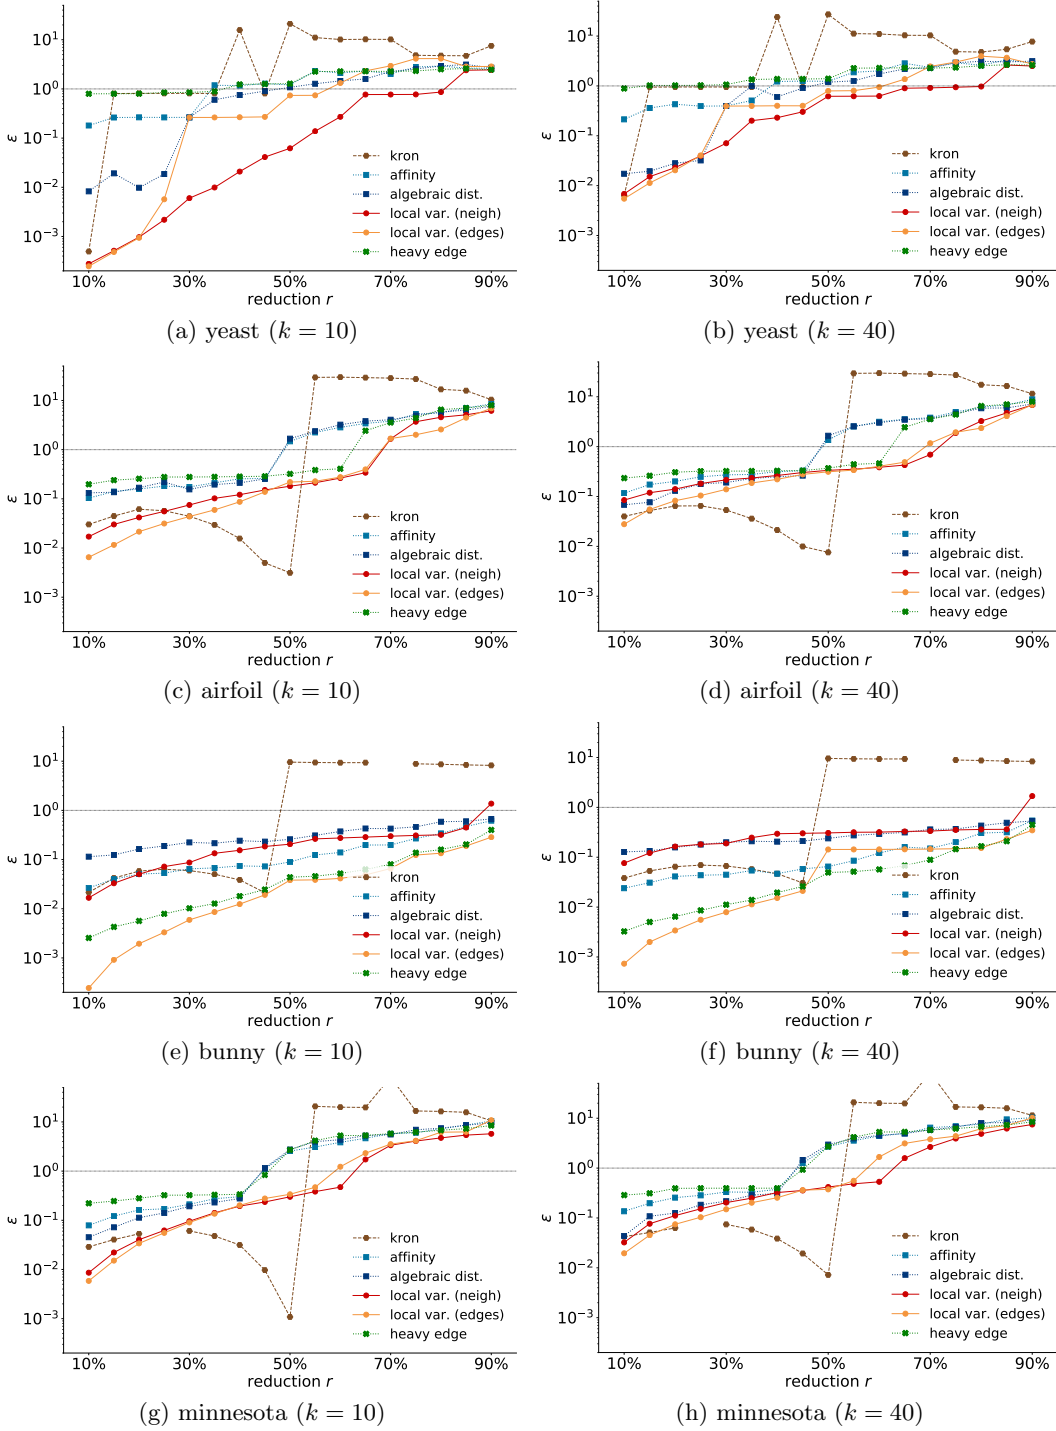


Figure 2: Quality of approximation comparison for four representative graphs (rows) and two subspace sizes (columns).

encouraging result, illustrating that significant dimensionality reduction is often possible, without sacrificing too much the solution quality.

Following intuition, it is generally harder to approximate subspaces of larger dimension k , but not excessively so. Increasing k from 10 to 40 in most cases increases ϵ only slightly. The only case where the approximation becomes profoundly better with small subspaces is with small reduction ratios r . For instance, coarsening the yeast graph results in an impressive approximation for all $r < 30\%$ when $k = 10$, whereas ϵ increases almost by an order of magnitude when k becomes 40.

Kron reduction is an effective way to half the graph size but can result in poor approximation otherwise. If one is willing to sacrifice in terms of efficiency, Kron reduction effectively reduces the graph size by a factor of two (with the exception of the yeast graph). What might be startling is that the method behaves poorly for different reduction ratios. Three main factors cause this deterioration of performance. First, the sampling set is constructed based on the sign pattern of u_N , i.e., by keeping the vertices v_i for which $u_N(i) \geq 0$, and has cardinality close to $N/2$ (Shuman et al., 2016). Therefore, if at any level one tries to reduce the graph size by less than half, the last eigenvector heuristic cannot be used precisely. Second, numerical instability issues sometimes manifest when r exceeds 50%. I was not able to improve the implementation featured in the PyGSP toolbox, and in my experiments some problem instances could not be solved successfully (hence the missing markers). The third reason is described next.

Coarse levels should aim to approximate the original graph and not the proceeding levels. The conventional approach in multi-level schemes is to aim at each level to approximate as closely as possible the graph of the previous level. This can lead to a sudden increase of error at consecutive levels (e.g. notice the minnesota error as r approaches 50%) as decisions early in the scheme can have a significant impact later on. On the other hand, local variation methods modify the cost function minimized at each level (see Proposition 17), resulting in smoother transitions between levels and tighter approximations at large r .

5.2. Spectrum Approximation

The second part of the experiments examines the coarsening through the lens of spectral graph theory. The premise is that, since the spectrum of the Laplacian distills information about the graph structure, one may interpret the spectral distance as a proxy for the structural similarity of the two graphs. This is by no means a new idea—the Laplacian spectrum is a common ingredient in accessing graph similarity (Wilson and Zhu, 2008).

Tables 1 and 2 report the mean relative eigenvalue error defined as $\frac{1}{k} \sum_{i=1}^k \frac{|\tilde{\lambda}_i - \lambda_i|}{\lambda_i}$ for two representative k , respectively 10 and 40. The results for $k = 80$ were consistent with those presented here, and are not reported here for reasons of brevity.

As expected, the reduction ratio plays a significant role in the closeness of Laplacian spectra. Indeed, for most cases, the eigenvalue error jumps by almost an order of magnitude whenever r increases by 20%. Yet, also in most cases, acceptable errors can be achieved even when

	r	heavy edge	local var. (edges)	local var. (neigh.)	algebraic distance	affinity	Kron reduction
yeast	30%	0.284	0.123	0.003	0.126	0.164	0.054
	50%	1.069	0.460	0.034	0.759	0.877	1.321
	70%	5.126	3.920	0.409	3.395	3.140	1.865
airfoil	30%	0.278	0.036	0.065	0.219	0.258	0.345
	50%	0.527	0.201	0.197	1.221	1.291	0.900
	70%	3.954	1.042	0.926	5.562	5.145	2.027
bunny	30%	0.015	0.006	0.061	0.244	0.070	0.335
	50%	0.064	0.046	0.190	0.401	0.137	0.801
	70%	0.122	0.080	0.323	0.694	0.304	1.812
minnesota	30%	0.332	0.088	0.078	0.220	0.295	0.324
	50%	1.363	0.431	0.310	2.394	2.676	0.873
	70%	7.452	4.553	1.892	8.412	9.354	2.068

Table 1: Mean relative error for the first $k = 10$ eigenvalues, for different graphs, reduction ratios, and coarsening methods.

	r	heavy edge	local var. (edges)	local var. (neigh.)	algebraic distance	affinity	Kron reduction
yeast	30%	0.311	0.113	0.023	0.113	0.162	0.120
	50%	1.087	0.413	0.130	0.522	0.676	1.196
	70%	3.618	2.212	0.454	2.387	2.474	1.946
airfoil	30%	0.278	0.095	0.181	0.191	0.257	0.366
	50%	0.555	0.326	0.349	0.692	0.788	0.955
	70%	2.059	0.905	0.848	2.309	2.619	2.141
bunny	30%	0.014	0.008	0.085	0.215	0.049	0.294
	50%	0.067	0.058	0.181	0.356	0.092	0.660
	70%	0.122	0.098	0.299	0.520	0.200	1.192
minnesota	30%	0.358	0.118	0.115	0.223	0.308	0.342
	50%	0.967	0.468	0.383	1.105	1.296	0.944
	70%	3.588	2.160	1.610	4.113	4.133	2.198

Table 2: Mean relative error for the first $k = 40$ eigenvalues, for different graphs, reduction ratios, and coarsening methods.

the coarse graph is as small as one third of the size of the original graph (corresponding to $r = 70\%$).

It might also be interesting to observe that there is a general agreement between the trends reported here and those described in the matrix approximation experiment. In particular, if one sorts the tested methods from best to worse, he/she will obtain an ordering that is generally consistent across the two experiments, with local variation methods giving the best approximation by a significant margin. A case in point: for the maximum ratio, the smallest local variation error is on average $3.5\times$ smaller than the leading state-of-the-art coarsening method. The gain is $2.3\times$ if the Kron reduction is also included in the comparison.

Overall, it can be deduced from these results that local variation methods coarsen a graph in a manner that preserves its spectrum. This is in accordance with the theoretical results. As it was shown by Theorem 13, if L and L_c act similar w.r.t. all vectors in \mathbf{U}_k , then their

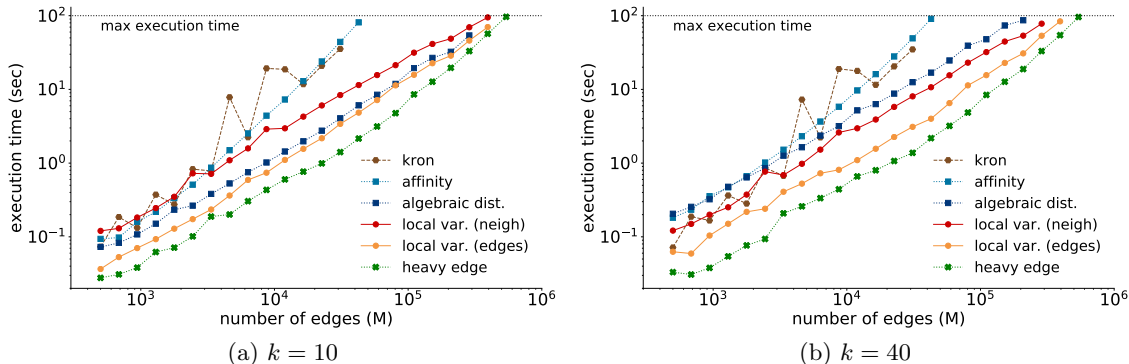


Figure 3: Execution time as a function of the number of edges of the graph.

eigenvalues cannot be far apart. Therefore, by aiming for restricted approximation, local variation methods implicitly also guarantee spectrum approximation.

5.3. Efficiency

The last experiment tests computational efficiency. I adopt a simple approach and aim to coarsen a 10-regular graph of increasing size⁸. I measure the execution time of the six different methods for graph reduction and report the mean over ten iterations while capping computation at 100 seconds.

The results are displayed in Figures 3a and 3b for subspaces of size 10 and 40, respectively. As with most such comparisons, the actual numbers are only indicative and depend on the programming language utilized (Python), processing paradigm (no parallelism was employed), and hardware architecture (2.2GHz CPU)⁹.

Focusing on the trends, except for Kron reduction and affinity, most methods scale quasi-linearly with the number of edges. Interestingly, local variation methods are quite competitive and do not sacrifice much as compared to the straightforward heavy edge matching. One can also observe that constructing $\mathcal{F}_{\ell-1}$ based on neighborhoods results in slightly slower computation than with the edge-based method construction because in the latter case the local variation cost can be computed much more efficiently.

6. Conclusions

Graph reduction techniques are commonly used in modern graph processing and machine learning pipelines in order to facilitate the processing and analysis of large graphs. Nevertheless, more often than not these techniques have been designed based on intuition and possess no rigorous justification or provable guarantees.

8. Starting from a ring, this graph is constructed by connecting each vertex with the five neighbors in either side.

9. I expect that a significant speedup can be achieved by compiling the code to native machine instructions as well as by parallelizing the local variation cost computation.

This work considered the graph reduction problem from the perspective of restricted spectral approximation—a modification of the measure utilized for the construction of graph sparsifiers. This measure is especially relevant when restricted to Laplacian eigenspaces of small variation, as it implies strong spectral and cut guarantees. The analysis of restricted spectral approximation has led to new algorithms for graph coarsening that provably approximate (a portion of) the graph spectrum, as well as the cut structure of a large graph.

Certain important questions remain open at the point of concluding this manuscript. To begin with, I am currently unaware of a rigorous way to determine how much one may benefit from reduction—that is, how small can ϵ be for a specific subspace and target n ? In addition, no polynomial-time algorithm for graph coarsening exists that provably approximates the minimal achievable ϵ . Finally, though I lack a formal proof, I also suspect that stronger cut guarantees can be derived from restricted spectral approximation. I argue that the potential of graph reduction cannot be fully realized until these fundamental questions are satisfactorily addressed.

Acknowledgments

I would like to thank the Swiss National Science Foundation (SNSF) for supporting this work in the context of the project “*Deep Learning for Graph-Structured Data*”, grant number PZ00P2 179981. I would like to thank Rodrigo C. G. Pena and Nguyen Q. Tran as well as the editor Karsten Borgwardt and the anonymous reviewers for their useful comments.

Appendix A. Deferred Proofs

A.1. Proof of Property 1

I draw up an inductive argument demonstrating that Π is a projection matrix. The *base case*, i.e., that $A_c = P_c^+ P_c$ is a projection matrix follows by the definition of the pseudo-inverse: $A_c A_c = P_c^+ P_c P_c^+ P_c = P_c^+ P_c = A_c$, where one uses the property $P_c = P_c P_c^+ P_c$.

For the *inductive step*, I argue that if $A_{\ell+1}$ is a projection matrix the same holds for

$$A_\ell = P_\ell^+ A_{\ell+1} P_\ell.$$

To this end, let $P_\ell = U \Sigma V^\top$ be the singular-value decomposition with $\Sigma = (D; 0) \in \mathbb{R}^{N_{\ell+1} \times N_\ell}$ decomposed into the $N_{\ell+1} \times N_{\ell+1}$ diagonal matrix D and the all zero matrix 0 . Then

$$A_\ell = V \Sigma^+ U^\top A_{\ell+1} U \Sigma V^\top.$$

Recalling that a projection matrix remains projective if it undergoes a similarity transformation, we deduce that $U^\top A_{\ell+1} U$ is also projective and, moreover, if $\Sigma^+ U^\top A_{\ell+1} U \Sigma$ is a projection matrix, so is A_ℓ . However, one may write

$$\Sigma^+ U^\top A_{\ell+1} U \Sigma = \begin{pmatrix} D^{-1} & 0 \\ 0 & 0 \end{pmatrix} U^\top A_{\ell+1} U \begin{pmatrix} D & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} D^{-1} U^\top A_{\ell+1} U D & 0 \\ 0 & 0 \end{pmatrix}.$$

As a block diagonal matrix whose blocks are projective (again $D^{-1}U^\top A_{\ell+1}UD$ is a similarity transformation), $\Sigma^+U^\top A_{\ell+1}U\Sigma$ is also a projection matrix. The proof that Π is a projection matrix concludes by letting the induction unfold backwards from c to 1 and setting $\Pi = A_1$.

A.2. Proof of Theorem 3

The Courant-Fischer min-max theorem for a Hermitian matrix L reads

$$\lambda_k = \min_{\dim(\mathbf{U})=k} \max_{x \in \mathbf{U}} \left\{ \frac{x^\top Lx}{x^\top x} \mid x \neq 0 \right\}, \quad (8)$$

whereas the same theorem for L_c gives

$$\tilde{\lambda}_k = \min_{\dim(\mathbf{U}_c)=k} \max_{x_c \in \mathbf{U}_c} \left\{ \frac{x_c^\top L_c x_c}{x_c^\top x_c} \mid x_c \neq 0 \right\} = \min_{\dim(\mathbf{U}_c)=k} \max_{Px \in \mathbf{U}_c} \left\{ \frac{x^\top \Pi L \Pi x}{x^\top P^\top P x} \mid x \neq 0 \right\},$$

where in the second equality I substitute $L_c = P^\top L P^+$ and $x_c = Px$.

We will need the following result:

Lemma 19 *For any P with full-row rank, the following holds:*

$$\lambda_1(P P^\top) x^\top \Pi x \leq x^\top P^\top P x \leq \lambda_n(P P^\top) x^\top \Pi x,$$

with $\lambda_1(P P^\top)$ and $\lambda_n(P P^\top)$, respectively the smallest and largest eigenvalues of $P P^\top$.

Proof Set $D = (P P^\top)^+$, which is an $n \times n$ PSD matrix. By the properties of the Moore–Penrose inverse $P^+ = P^\top (P P^\top)^+ = P^\top D$ and therefore $P^\top P = P^\top D D^{-1} P = P^+ D^{-1} P$. Supposing that the eigenvalues of D lie in $[\lambda_{\min}(D), \lambda_{\max}(D)]$ and that P is full row-rank such that $\lambda_{\min}(D) > 0$, one may write

$$\frac{1}{\lambda_{\max}(D)} x^\top \Pi x \leq x^\top P^\top P x = x^\top P^+ D^{-1} P x \leq \frac{1}{\lambda_{\min}(D)} x^\top \Pi x.$$

To grasp why the aforementioned inequality holds, first use the cyclic property of the trace to obtain

$$x^\top P^+ D^{-1} P x = \text{tr}(x^\top P^+ D^{-1} P x) = \text{tr}(D^{-1} P x x^\top P^+),$$

and further recall that for any symmetric A and PSD matrix B one has $\lambda_{\min}(A)\text{tr}(B) \leq \text{tr}(AB) \leq \lambda_{\max}(A)\text{tr}(B)$, where $\lambda_{\min}(A)$ denotes the smallest eigenvalue of A (and $\lambda_{\max}(A)$ is the largest) (Fang et al., 1994). Set $A = D^{-1}$, which is by assumption symmetric, and $B = P x x^\top P^+$, which is PSD since its rank is (at most) one with the only (potentially) non-zero eigenvalue exactly $\text{tr}(P x x^\top P^+) = x^\top P^+ P x = x^\top \Pi x = x^\top \Pi^2 x = \|\Pi x\|_2^2 \geq 0$, where $\Pi = \Pi^2$ due to being a projection matrix (Property 1). The desired inequality then follows since $\lambda_{\min}(D^{-1}) = 1/\lambda_{\max}(D) = \frac{1}{\beta}$, $\lambda_{\max}(D^{-1}) = 1/\lambda_{\min}(D) = \frac{1}{\alpha}$ and once more $\text{tr}(P x x^\top P^+) = x^\top \Pi x$.

Finally, since P is full row-rank, D is invertible meaning $\lambda_{\min}(D) = 1/\lambda_n(PP^\top)$ and $\lambda_{\max}(D) = 1/\lambda_1(PP^\top)$. \blacksquare

From the above, it is deduced that

$$\begin{aligned} \tilde{\lambda}_k &\geq \min_{\dim(\mathbf{U}_c)=k} \max_{P x \in \mathbf{U}_c} \left\{ \frac{x^\top \Pi L \Pi x}{\lambda_n(PP^\top) x^\top \Pi x} \mid x \neq 0 \right\} \\ &= \frac{1}{\lambda_n(PP^\top)} \min_{\dim(\mathbf{U})=k, \mathbf{U} \subseteq \text{im}(\Pi)} \max_{x \in \mathbf{U}} \left\{ \frac{x^\top L x}{x^\top x} \mid x \neq 0 \right\}, \end{aligned}$$

where the equality holds since Π is a projection matrix (see Property 1). Notice how, with the exception of the constraint $x = \Pi x$ and the multiplicative term, the final optimization problem is identical to the one for λ_k , given in (8). As such, the former's solution must be strictly larger (since it is a more constrained problem) and $\tilde{\lambda}_k \geq \frac{\lambda_k}{\lambda_n(PP^\top)}$.

Analogously, one obtains the lower inequality $\tilde{\lambda}_{k-(N-n)} \leq \frac{\lambda_k}{\lambda_1(PP^\top)}$ by applying the same argument on matrices $-L$ and $-L_c$ and exploiting that the k -th largest eigenvalue of any matrix M is also the k -th smallest eigenvalue of $-M$.

A.3. Proof of Proposition 6

For notational convenience, I drop the level index supposing that $c = 1$ and thus P is an $n \times N$ coarsening matrix. As we will see in the following, P has rank n and thus to prove that $P^+ = P^\top D^{-2}$, it is sufficient to show that matrix $\Pi = P^\top D^{-2} P$ is a projection matrix of rank n (and thus equal to $P^+ P$). Matrix $\tilde{P} = D^{-1} P$ has the same sparsity structure as P and is thus also a coarsening matrix. W.l.o.g. let the rows of P be sorted based on their support, such that for any two rows $r < r'$ and $P(r, i), P(r', i') \neq 0$ we necessarily have $i < i'$. Furthermore, denote by p_r the vector containing all non-zero entries of $P(r, :)$ such that $\|p_r\|_2 = \|P(r, :)\|_2 = D(r, r)$. Due to the disjoint support of the rows of P and under this particular sorting, matrix Π is block-diagonal. Moreover, each block B_r in its diagonal is a rank 1 projection matrix as $B_r^2 = B_r B_r = (p_r D(r, r)^{-2} p_r^\top) (p_r D(r, r)^{-2} p_r^\top) = p_r D(r, r)^{-2} p_r^\top \frac{p_r p_r^\top}{\|p_r\|_2^2} = B_r$. We have thus arrived to the relation $\Pi^2 = \Pi$, which constitutes a necessary and sufficient condition for Π to be a projection matrix. The block-diagonal structure of Π also implies that its rank (as well as that of P) is n .

A.4. Proof of Proposition 7

Let us first remark that, by Proposition 6, $A = P^\mp$ is also a coarsening matrix with the same sparsity structure as P .

Necessity. I start by considering the nullspace of $L_c = A L A^\top$ and aim to ensure that it is equal to the span of the constant vector $\mathbf{1}$, which is a necessity for all combinatorial Laplacian matrices. Since matrix A is full row-rank and L has rank $N - 1$, the nullspace of L_c is one dimensional. Therefore, the nullspace is correct as long as $\mathbf{1}^\top A L A^\top \mathbf{1} = 0$, which

happens if either $A^\top \mathbf{1} = \alpha \mathbf{1}$ for a constant α or $A^\top \mathbf{1} = 0$. In both cases, $(A^\top \mathbf{1})(r) = \alpha \mathbf{1}$ for every r . By the definition of A however, we know that its rows have disjoint support and, as such, vector $A^\top \mathbf{1}$ exactly contains the non-zero entries of A . In other words, for the nullspace of L_c to be properly formed, the non-zero entries of A should either all be equal to α (such that $A^\top \mathbf{1} = \alpha \mathbf{1}$) or zero (in which case $A^\top \mathbf{1} = 0$). The latter case can clearly be discarded as it would disconnect the graph. We have thus discovered that ALA^\top has a properly formed nullspace if and only if the non-zero entries of A are equal, rendering the latter condition necessary.

Sufficiency. Considering that every Laplacian of M edges can be re-written as $L = S^\top S$, where S is the $M \times N$ incidence matrix, one may confirm that the condition is also sufficient by showing that, for every A with equal non-zero entries, the matrix $S_c = SA^\top$ is an incidence matrix of L_c such that $L_c = S_c^\top S_c$. W.l.o.g., suppose that $\alpha = 1$ ($\alpha^2 L$ is a valid Laplacian for all α). By construction, each row of S_c is $S_c(q, :)^\top = AS(q, :)^\top$. Name as e_{ij} the corresponding edge, such that $S(q, :)^\top = \delta_i - \delta_j$, where δ_i is a dirac centered at vertex v_i . It follows that $S_c(q, :)^\top = A\delta_i - A\delta_j$. Obviously, if none of v_i, v_j are contracted then $S_c(q, :)^\top = \delta_i - \delta_j$, which is a valid row. Moreover, by construction of A , if either of v_i, v_j is contracted (but not both) or if v_i, v_j are contracted into different vertices then both relations $A\delta_i = \delta_i$ and $A\delta_j = \delta_j$ hold and thus once more $S_c(q, :)^\top = \delta_i - \delta_j$ is a valid row. Lastly, if v_i, v_j are contracted into the same vertex then for some r it must be that $A(r, i) = A(r, j)$, whereas $A(r', i) = A(r', j) = 0$ for all $r' \neq r$ and thus $S_c(r, :)^\top = 0$, signifying that the edge is not present. Summarizing, in every case S_c is a valid incidence matrix, rendering the condition also sufficient.

A.5. Proof of Property 8

For any two disjoint subsets $\mathcal{S}_1, \mathcal{S}_2$ of \mathcal{V} denote by $w(\mathcal{S}_1, \mathcal{S}_2) = \sum_{v_i \in \mathcal{S}_1} \sum_{v_j \in \mathcal{S}_2} w_{ij}$ the cut weight in G .

The claim is proven by induction on the number of levels. For the *base case*, set $\ell = 1$ and define $C = P_1^+$ such that $c_r = C(:, r)$ is the indicator vector of the contraction set $\mathcal{V}_0^{(r)}$. It is a consequence of the Laplacian form of L that, for any $v_r, v_q \in \mathcal{V}_1$ with $r \neq q$, we have

$$\begin{aligned} W_1(r, q) &= -L_1(r, q) = -c_r^\top L c_q = \sum_{v_i \neq v_j} w_{ij} c_r(i) c_q(j) + \sum_{v_i} d_i c_r(i) c_q(i) \\ &= \sum_{v_i \in \mathcal{V}_0^{(r)}, v_j \in \mathcal{V}_0^{(q)}} w_{ij} = w(\mathcal{S}_1^{(r)}, \mathcal{S}_1^{(q)}), \end{aligned}$$

where the penultimate step uses $c_r(i) c_q(i) = 0$ since contraction steps are disjoint, and the last step exploits the equivalence $\mathcal{V}_0^{(q)} = \mathcal{S}_1^{(q)}$. For the *inductive step*, consider level $\ell > 1$. Since $L_{\ell-1}$ is a Laplacian matrix, one may employ an identical argument as when $\ell = 1$ to find that the weight between vertices $v_r, v_q \in \mathcal{V}_\ell$ with $r \neq q$ is

$$W_\ell(r, q) = \sum_{v_i \in \mathcal{V}_{\ell-1}^{(r)}, v_j \in \mathcal{V}_{\ell-1}^{(q)}} W_{\ell-1}(i, j).$$

By the induction hypothesis however, it must be $W_{\ell-1}(i, j) = w(\mathcal{S}_{\ell-1}^{(i)}, \mathcal{S}_{\ell-1}^{(j)})$, implying

$$W_{\ell}(r, q) = \sum_{v_i \in \mathcal{V}_{\ell-1}^{(r)}, v_j \in \mathcal{V}_{\ell-1}^{(q)}} w(\mathcal{S}_{\ell-1}^{(i)}, \mathcal{S}_{\ell-1}^{(j)}) = w(\mathcal{S}_{\ell}^{(r)}, \mathcal{S}_{\ell}^{(q)}),$$

with the final equality being true due to the recursive definition of sets $\mathcal{S}_{\ell}^{(r)}$ and $\mathcal{S}_{\ell}^{(q)}$, as well as the following property of cuts: for any two sets (call them large sets) and any partition of each into an arbitrary number of subsets, the cut between the large sets is equal to the sum of all cuts between pairs of subsets belonging to different large sets. This completes the proof.

A.6. Proof of Theorem 13

The lower bound is given by Theorem 3. For the upper bound, I reason similarly to the proof of the latter to find:

$$\tilde{\lambda}_k = \min_{\dim(\mathbf{U}_c)=k} \max_{x_c \in \mathbf{U}_c} \left\{ \frac{x_c^\top L_c x_c}{x_c^\top x_c} \mid x_c \neq 0 \right\} \leq \gamma_2 \min_{\dim(\mathbf{U})=k, \mathbf{U} \subseteq \text{im}(\Pi)} \max_{x \in \mathbf{U}} \left\{ \frac{x^\top \Pi L \Pi x}{x^\top \Pi x} \mid x \neq 0 \right\}.$$

Above, the inequality is due to Lemma 19 with $\gamma_2 = 1/\lambda_1(PP^\top)$. Thus, for any matrix V the following inequality holds

$$\tilde{\lambda}_k \leq \gamma_2 \max_{x \in \text{span}(V)} \left\{ \frac{x^\top \Pi L \Pi x}{x^\top \Pi x} \mid x \neq 0 \right\},$$

as long as the image of V is of dimension k and does not intersect the nullspace of Π . Write U_k to denote the $N \times k$ matrix with the k first eigenvectors of L , whose image is of dimension k as needed. Assume for now that the nullspace requirement is also met:

$$\tilde{\lambda}_k \leq \gamma_2 \max_{x \in \text{span}(U_k)} \left\{ \frac{x^\top \Pi L \Pi x}{x^\top \Pi x} \mid x \neq 0 \right\} = \gamma_2 \max_{x \in \text{span}(U_k)} \left\{ \frac{\|S\Pi x\|_2^2}{\|\Pi x\|_2^2} \mid x \neq 0 \right\}.$$

It will be convenient to manipulate the square-root of this quantity:

$$\sqrt{\frac{\tilde{\lambda}_k}{\gamma_2}} \leq \max_{a \in \mathbb{R}^k} \frac{\|S\Pi U_k a\|_2}{\|\Pi U_k a\|_2} = \frac{\|S\Pi U_k\|_2}{\|\Pi U_k\|_2} \leq \frac{\|S U_k\|_2 + \|S\Pi^\perp U_k\|_2}{\|\Pi U_k\|_2} = \frac{\sqrt{\lambda_k} + \|S\Pi^\perp U_k\|_2}{\|\Pi U_k\|_2}, \quad (9)$$

with S defined such that $L = S^\top S$. The norm in the numerator is upper bounded by

$$\begin{aligned} \|S\Pi^\perp U_k\|_2 &= \|S\Pi^\perp U_k \Lambda_k^{-1/2} \Lambda_k^{1/2}\|_2 \leq \|S\Pi^\perp U_k \Lambda_k^{+1/2}\|_2 \|\Lambda_k^{1/2}\|_2 \\ &= \sqrt{\lambda_k} \|S\Pi^\perp U_k \Lambda_k^{+1/2}\|_2 = \sqrt{\lambda_k} \epsilon_k. \end{aligned}$$

If the last step is not immediately obvious, one can be convinced by first exploiting the unitary-invariance of the spectral norm to write $\|S\Pi^\perp U_k \Lambda_k^{+1/2}\|_2 = \|S\Pi^\perp U_k U_k^\top L^{+1/2}\|_2$, and then confirming in the proof of Proposition 17 that the latter quantity is exactly ϵ_k when $V = U_k$.

The preceding analysis assumed that the image of U_k and the nullspace of Π did not intersect. Since $\Pi^\perp = I - \Pi$ is a complement projection matrix, the previous holds when $\|\Pi^\perp U_k\|_2^2 < 1$. Since $\|\Pi^\perp u_1\| = 0$, one may w.l.o.g. exclude the first eigenvector u_1 from the space of interest. For the remainder of $\text{im}(U_k)$ the following holds:

$$\|\Pi^\perp U_k\|_2^2 = \max_{x \in \mathbf{U}_k \text{ and } x \perp u_1} \frac{\|\Pi^\perp x\|_2^2}{\|x\|_2^2} \leq \frac{1}{\lambda_2} \max_{x \in \mathbf{U}_k \text{ and } x \perp u_1} \frac{\|\Pi^\perp x\|_L^2}{\|x\|_2^2} = \epsilon_k^2 \frac{\lambda_k}{\lambda_2}.$$

Therefore, when $\epsilon_k^2 < \frac{\lambda_2}{\lambda_k}$, the nullspace condition is met. The proof is then concluded by substituting the bound $\|\Pi U_k\|_2^2 = 1 - \|\Pi^\perp U_k\|_2^2 \geq 1 - \epsilon_k^2 \frac{\lambda_k}{\lambda_2}$ in the denominator of (9).

A.7. Proof of Theorem 14

Li (1994) showed that we can express the $\sin \Theta$ as a sum of squared inner products:

$$\left\| \sin \Theta(U_k, P^\top \tilde{U}_k) \right\|_F^2 = \left\| \tilde{U}_{k^\perp}^\top P U_k \right\|_F^2 = \sum_{i \leq k} \sum_{j > k} (\tilde{u}_j^\top P u_i)^2 \quad (10)$$

If L_c and L are $(\mathbf{U}_i, \epsilon_i)$ -similar it follows from Corollary 12 that

$$u_i^\top P^\top L_c P u_i \leq (1 + \epsilon_i)^2 \lambda_i.$$

Summing these inequalities for all $i \leq k$ amounts to

$$\begin{aligned} \sum_{i \leq k} (1 + \epsilon_i)^2 \lambda_i &\geq \sum_{i \leq k} \sum_{j=1}^n \tilde{\lambda}_j (\tilde{u}_j^\top P u_i)^2 \\ &\geq \gamma_1 \sum_{i \leq k} \sum_{j=1}^n \lambda_j (\tilde{u}_j^\top P u_i)^2 \\ &= \gamma_1 \sum_{j \leq k} \lambda_j \sum_{i \leq k} (\tilde{u}_j^\top P u_i)^2 + \gamma_1 \sum_{j > k} \lambda_j \sum_{i \leq k} (\tilde{u}_j^\top P u_i)^2. \end{aligned} \quad (11)$$

where, following Theorem 3, I set $\gamma_1 = 1/\lambda_n(P P^\top)$ such that $\tilde{\lambda}_i \geq \gamma_1 \lambda_i$. Perform the following manipulation:

$$\begin{aligned} \sum_{j \leq k} \lambda_j \sum_{i \leq k} (\tilde{u}_j^\top P u_i)^2 &\geq \sum_{j \leq k} \lambda_j \sum_{i \leq k} (\tilde{u}_j^\top P u_i)^2 = \sum_{j \leq k} \lambda_j \left(1 - \sum_{i > k} (\tilde{u}_j^\top P u_i)^2 \right) \\ &\geq \sum_{j \leq k} \lambda_j - \lambda_k \sum_{i \leq k} \left(\|\Pi^\perp u_i\|_2^2 + \sum_{j \geq k} (\tilde{u}_j^\top P u_i)^2 \right), \end{aligned}$$

which together with (10) and (11) yields

$$\left\| \sin \Theta(U_k, P^\top \tilde{U}_k) \right\|_F^2 \leq \sum_{i \leq k} \frac{(1 + \epsilon_i)^2 \lambda_i / \gamma_1 - \lambda_i + \lambda_k \|\Pi^\perp u_i\|_2^2}{\lambda_{k+1} - \lambda_k} \quad (12)$$

To proceed, I note the following useful inequality:

Lemma 20 *If L and L_c are $(\mathbf{U}_i, \epsilon_i)$ -similar, then $\|\Pi^\perp u_i\|_2^2 \leq \epsilon_i$.*

Proof For every u_i one sees that

$$\begin{aligned} u_i^\top P^\top L_c P u_i &= u_i^\top \Pi L \Pi u_i = u_i^\top (I - \Pi^\perp) L (I - \Pi^\perp) u_i \\ &= u_i^\top L u_i - 2u_i^\top L \Pi^\perp u_i + u_i^\top \Pi^\perp L \Pi^\perp u_i \\ &= \lambda_i - 2\lambda_i u_i^\top \Pi^\perp u_i + u_i^\top \Pi^\perp L \Pi^\perp u_i \end{aligned}$$

meaning that

$$\|\Pi^\perp u_i\|_2^2 = \frac{1}{2} \left(1 + \frac{u_i^\top \Pi^\perp L \Pi^\perp u_i}{\lambda_i} - \frac{u_i^\top P^\top L_c P u_i}{\lambda_i} \right) \leq \frac{1}{2} (1 + \epsilon_i^2 - (1 - \epsilon_i)^2) = \epsilon_i.$$

The last inequality is because, by restricted spectral approximation, we have $u_i^\top \Pi^\perp L \Pi^\perp u_i = \|\Pi^\perp u_i\|_L^2 \leq \epsilon_i^2 \|u_i\|_L^2 = \epsilon_i^2 \lambda_i$ and from Corollary 12:

$$u_i^\top P^\top L_c P u_i = \|P u_i\|_{L_c}^2 \geq (1 - \epsilon_i)^2 \|u_i\|_L^2 = (1 - \epsilon_i)^2 \lambda_i. \quad \blacksquare$$

As a consequence, it follows that

$$\left\| \sin \Theta(U_k, P^\top \tilde{U}_k) \right\|_F^2 \leq \sum_{i \leq k} \frac{(1 + \epsilon_i)^2 \lambda_i / \gamma_1 - \lambda_i + \lambda_k \epsilon_i}{\lambda_{k+1} - \lambda_k},$$

which after manipulation gives the desired inequality.

A.8. Proof of Theorem 15

The lower bound is a direct consequence of consistent coarsening and holds independently of restricted spectral approximation: for any set $\mathcal{S}_c \subset \mathcal{V}_c$ define $\mathcal{S} \subset \mathcal{V}$ such that $v_i \in \mathcal{S}$ if and only if $\varphi_c \circ \varphi_{c-1} \circ \dots \circ \varphi_1(v_i) \in \mathcal{S}_c$. Clearly, $w(\mathcal{S}) \geq w_c(\mathcal{S}_c)$, where the subscript c implies that the latter volume is w.r.t. G_c . In addition, by the definition of Laplacian consistent coarsening and since every contraction set belongs either in \mathcal{S} or $\bar{\mathcal{S}}$ (but not in both), it follows that $w(\mathcal{S}, \bar{\mathcal{S}}) = w_c(\mathcal{S}_c, \bar{\mathcal{S}}_c)$. In other words, for every \mathcal{S}_c there exists a set \mathcal{S} such that

$$\phi(\mathcal{S}) = \frac{w(\mathcal{S}, \bar{\mathcal{S}})}{\min\{w(\mathcal{S}), w(\bar{\mathcal{S}})\}} \leq \frac{w_c(\mathcal{S}_c, \bar{\mathcal{S}}_c)}{\min\{w_c(\mathcal{S}_c), w_c(\bar{\mathcal{S}}_c)\}} = \phi_c(\mathcal{S}_c),$$

implying also that the k -conductance of G and G_c are related by $\phi_k(G) \leq \phi_k(G_c)$.

For the upper bound, I exploit the following multi-way Cheeger inequality:

Theorem 21 (Restatement of Theorem 1.2 by Lee et al. (2014)) *For every graph G and every $k \in \mathbb{N}$, we have*

$$\frac{\mu_k}{2} \leq \phi_k(G) = O(\sqrt{\mu_{2k} \xi_k(G)}),$$

with $\xi_k(G) \leq \log k$. If G is planar then $\xi_k(G) \leq 1$ and if G excludes K_h as a minor (but is not planar) then $\xi_k(G) \leq h^4$.

Further, in the standard Cheeger inequality (Alon and Milman, 1985; Alon, 1986) $k = 2$ and the upper bound is given by $\sqrt{2\mu_2}$. Note that the eigenvalues mentioned here are those of the normalized Laplacian matrix $L^n = D^{-1/2}LD^{-1/2}$. To this end, suppose that V_{2k} contains the first $2k$ eigenvectors of L^n and fix $\mathbf{R} = \text{span}(D^{-1/2}V_{2k})$. Perform consistent coarsening w.r.t. to the combinatorial Laplacian L (not L^n). Then, by definition, if L_c and L are $(\epsilon_{2k}, \mathbf{R})$ -similar then for every $x \in \mathbf{R}$ one gets

$$\|\Pi^\perp x\|_L \leq \epsilon_{2k} \|x\|_L.$$

The substitution $y = D^{1/2}x$, such that $y \in \text{span}(V_{2k})$, allows us to transform the semi-norms above into semi-norms concerning L^n as follows:

$$\|x\|_L^2 = x^\top Lx = y^\top D^{-1/2}LD^{-1/2}y = \|y\|_{L^n}^2$$

and

$$\|\Pi^\perp x\|_L = \|D^{-1/2}D^{1/2}\Pi^\perp D^{-1/2}D^{1/2}x\|_L = \|D^{1/2}\Pi^\perp D^{-1/2}y\|_{L^n} = \|(I - \Pi^n)y\|_{L^n}.$$

Above, $\Pi^n = D^{1/2}\Pi D^{-1/2}$ is the projection matrix (the set of projection matrices is closed under similarity transformations) corresponding to the coarsening matrix $P^n = PD^{-1/2}$, and now $y_c = \Pi^n y$. It follows that, for every $y \in \mathbf{V}_{2k} = \text{span}(V_{2k})$, we have

$$\|y - \Pi^n y\|_{L^n} \leq \epsilon_{2k} \|y\|_{L^n}$$

and thus L_c^n and L^n are $(\mathbf{V}_{2k}, \epsilon_{2k})$ -similar.

Combining the multi-way Cheeger inequality with Theorem 13 for L_c^n and L^n one obtains

$$\begin{aligned} \phi_k^2(G_c) &= O(\tilde{\mu}_{2k} \xi_k(G)) \\ &= O\left(\gamma_2(1 + \epsilon_{2k})^2 \frac{\mu_2}{\mu_2 - \epsilon_{2k}\mu_{2k}} \mu_{2k} \xi_k(G)\right) \\ &= O\left(\frac{\gamma_2(1 + \epsilon_{2k})^2 \xi_k(G)}{1 - \epsilon_{2k}^2(\mu_{2k}/\mu_2)} \phi_k(G)\right), \end{aligned}$$

where the eigenvalues above are those of L^n and the preceding holds whenever $\epsilon_{2k}^2 < \mu_2/\mu_{2k}$. Further, when $k = 2$ the upper bound simplifies to $\phi_2^2(G_c) \leq \frac{4\gamma_2(1+\epsilon_2)^2}{1-\epsilon_2^2} \phi_2(G)$.

A.9. Proof of Proposition 17

The following analysis is slightly more general than what is claimed in the statement of Proposition 17: it holds for arbitrary PSD L and L_c (i.e., not necessarily Laplacian matrices) as long as the image $\text{im}(\Pi)$ of the projection matrix Π encloses the nullspace of L . The former trivially holds for Laplacian consistent coarsening, as, by design, one has $\Pi\mathbf{1} = \mathbf{1}$ (see Section 2.4).

Let $V \in \mathbb{R}^{N \times k}$ be a basis of \mathbf{R} . I start by proving that, for any integer $k \leq n$ and for all $x \in \text{span}(V)$ the inequality

$$\|x - \Pi x\|_L \leq \epsilon \|x\|_L$$

holds for all $\epsilon \geq \|\Pi^\perp B_0\|_L$, where $B_0 = VV^\top L^{+1/2}$. I remind the reader that $\|x\|_L = \|Sx\|_2 = \|L^{1/2}x\|_2$ and $\Pi^\perp = I - P^+P$. Furthermore, since $\text{im}(\Pi)$ necessarily encloses the nullspace \mathbf{N} of L , w.l.o.g., one may assume that $\forall x \in \mathbf{R}$ the action of matrix L is invertible. To see why, note that if $x \in \mathbf{N}$ then $\|x\|_L = 0$ and $\|x - \Pi x\|_L = 0$, meaning that the inequality above is trivially satisfied. I then derive

$$\begin{aligned} \max_{x \in \mathbf{R}} \frac{\|x - \Pi x\|_L}{\|x\|_L} &= \max_{x \in \mathbf{R}} \frac{\|S\Pi^\perp x\|_2}{\|L^{1/2}x\|_2} \\ &= \max_{x \in \mathbf{R}} \frac{\|S\Pi^\perp VV^\top x\|_2}{\|L^{1/2}x\|_2} \end{aligned} \quad (13)$$

$$\begin{aligned} &= \max_{x \in \text{im}(LV)} \frac{\|S\Pi^\perp VV^\top L^{+1/2}x\|_2}{\|x\|_2} \\ &\leq \|S\Pi^\perp VV^\top L^{+1/2}\|_2 = \|\Pi^\perp B_0\|_L, \end{aligned} \quad (14)$$

where equality (13) holds because VV^\top is a projection onto \mathbf{R} , whereas equality (14) is true since L is invertible within \mathbf{R} .

One should also note that, for the specific case where V is an eigenspace of L , $\text{im}(LV) = \mathbf{R}$ and as such $\epsilon = \|\Pi^\perp x\|_L / \|x\|_L = \|S\Pi^\perp VV^\top L^{+1/2}\|_2$ (once more w.l.o.g. the nullspace of L can be ignored).

In addition, as the following technical lemma claims, in a multi-level scheme, any $\|\Pi^\perp x\|_L$ can be broken down into the contributions of each level:

Lemma 22 *Define projection matrices $\Pi_\ell = P_\ell^+ P_\ell$ and $\Pi_\ell^\perp = I - \Pi_\ell$. If*

$$\|\Pi_\ell^\perp x_{\ell-1}\|_{L_{\ell-1}} \leq \sigma_\ell \|x_{\ell-1}\|_{L_{\ell-1}} \quad \text{at each level } \ell \leq c,$$

then the multi-level error is bounded by

$$\|\Pi^\perp x\|_L \leq \left(\sum_{\ell=1}^c \sigma_\ell \prod_{q=1}^{\ell-1} (1 + \sigma_q) \right) \|x\|_L = \left(\prod_{\ell=1}^c (1 + \sigma_\ell) - 1 \right) \|x\|_L.$$

Proof Recursively apply the following inequality

$$\begin{aligned} \left\| S_{\ell-1} \Pi_\ell^\perp x_{\ell-1} \right\|_2 &\leq \sigma_\ell \|S_{\ell-1} x_{\ell-1}\|_2 \\ &= \sigma_\ell \|S_{\ell-2} \Pi_{\ell-1} x_{\ell-2}\|_2 \\ &\leq \sigma_\ell \left(\|S_{\ell-2} x_{\ell-2}\|_2 + \left\| S_{\ell-2} \Pi_{\ell-1}^\perp x_{\ell-2} \right\|_2 \right) \\ &\leq \sigma_\ell (\|S_{\ell-2} x_{\ell-2}\|_2 + \sigma_{\ell-1} \|S_{\ell-2} x_{\ell-2}\|_2) = \sigma_\ell (1 + \sigma_{\ell-1}) \|S_{\ell-2} x_{\ell-2}\|_2 \end{aligned}$$

to deduce that

$$\left\| S_{\ell-1} \Pi_\ell^\perp x_{\ell-1} \right\|_2 \leq \sigma_\ell \prod_{q=1}^{\ell-1} (1 + \sigma_q) \|S_0 x_0\|_2 = \sigma_\ell \prod_{q=1}^{\ell-1} (1 + \sigma_q) \|x\|_L.$$

The end-to-end error $\|S\Pi^\perp x\|_2$ is controlled with a simple telescopic series argument.

$$\begin{aligned} \|\Pi^\perp x\|_L &= \|S_0\Pi^\perp x_0\|_2 = \|S_0x_0 - S_c x_c\|_2 \\ &\leq \|S_0x_0 - S_1x_1\|_2 + \|S_1x_1 - S_2x_2\|_2 + \dots + \|S_{c-1}x_{c-1} - S_c x_c\|_2 \\ &= \|S_0\Pi_1^\perp x_1\|_2 + \|S_1\Pi_1^\perp x_1\|_2 + \dots + \|S_{c-1}\Pi_c^\perp x_{c-1}\|_2 \end{aligned}$$

Together, the above two results imply the desired bound. \blacksquare

Therefore, to guarantee that in a multi-level scheme

$$\|\Pi^\perp B_0\|_L = \max_{b \in \mathbb{R}^N} \frac{\|S\Pi^\perp B_0 b\|_2}{\|b\|_2} \leq \epsilon,$$

one needs to make sure that, for each level $\ell = 1, \dots, c$, the following holds:

$$\frac{\|S_{\ell-1}\Pi_\ell^\perp x_{\ell-1}\|_2}{\|S_{\ell-1}x_{\ell-1}\|_2} \leq \sigma_\ell, \quad \text{for all } x_{\ell-1} = P_{\ell-1} \cdots P_1 B_0 b$$

By the same argument used for the multi-level error, when $\ell = 1$, we have that $\sigma_1 = \|\Pi_1^\perp B_0\|_{L_0}$. For all other ℓ , set $B_{\ell-1} = P_{\ell-1} \cdots P_1 B_0$ and further let $(B_{\ell-1}^\top L_{\ell-1} B_{\ell-1})^{+1/2}$ be the pseudo-inverse of the matrix square-root of the $N \times N$ matrix $B_{\ell-1}^\top L_{\ell-1} B_{\ell-1}$. By the substitution $b = S_{\ell-1} B_{\ell-1} a$, the above can be rewritten as

$$\max_{b \in \mathbb{R}^N} \frac{\|S_{\ell-1}\Pi_\ell^\perp B_{\ell-1} b\|_2}{\|S_{\ell-1} B_{\ell-1} b\|_2} = \max_{b \in \mathbb{R}^N} \frac{\|S_{\ell-1}\Pi_\ell^\perp B_{\ell-1} (B_{\ell-1}^\top L_{\ell-1} B_{\ell-1})^{+1/2} b\|_2}{\|b\|_2}.$$

For $\ell > 1$, therefore $\sigma_\ell = \|\Pi_\ell^\perp A_{\ell-1}\|_{L_{\ell-1}}$ with $A_{\ell-1} = B_{\ell-1} (B_{\ell-1}^\top L_{\ell-1} B_{\ell-1})^{+1/2}$.

A.10. Proof of Proposition 18

For notational simplicity in the context of this proof I drop level indices and assume that only a single coarsening level is used—this is without loss of generality, as an identical argument holds for every level of the scheme.

Consider any x and set $y = \Pi^\perp x$. Furthermore, define for each contraction set the (i) *internal* edge set $\mathcal{E}^{(r)} = \{e_{ij} | v_i, v_j \in \mathcal{V}^{(r)} \text{ and } e_{ij} \in \mathcal{E}_{\ell-1}\}$, and (ii) the *boundary* edge set $\partial\mathcal{E}^{(r)}$, such that $e_{ij} \in \partial\mathcal{E}^{(r)}$ if and only if $v_i \in \mathcal{V}^{(r)}$ and $v_j \notin \mathcal{V}^{(r)}$. It is true that

$$\begin{aligned} \|\Pi^\perp x\|_L^2 &= \sum_{e_{ij} \in \mathcal{E}} w_{ij} (y(i) - y(j))^2 \\ &= \sum_{r=1}^n \left(\underbrace{\sum_{e_{ij} \in \mathcal{E}^{(r)}} w_{ij} (y(i) - y(j))^2}_{a_r} + \frac{1}{2} \underbrace{\sum_{e_{ij} \in \partial\mathcal{E}^{(r)}} w_{ij} (y(i) - y(j))^2}_{b_r} \right). \end{aligned}$$

In the following, I will express a_r and b_r as a function of the vector $y_r = \Pi_{\mathcal{V}(r)}^\perp x$. Term a_r is luckily independent of any other contraction set:

$$a_r = \sum_{e_{ij} \in \mathcal{E}(r)} w_{ij} (y(i) - y(j))^2 = \sum_{e_{ij} \in \mathcal{E}(r)} w_{ij} (y_r(i) - y_r(j))^2.$$

On the other hand, b_r is smaller than

$$b_r = \sum_{e_{ij} \in \partial \mathcal{E}(r)} w_{ij} (y(i) - y(j))^2 \leq 2 \sum_{e_{ij} \in \partial \mathcal{E}(r)} w_{ij} (y(i) - 0)^2 + 2 \sum_{e_{ij} \in \partial \mathcal{E}(r)} w_{ij} (0 - y(j))^2.$$

Distributing the second quantities, respectively, amongst the contraction sets that include said vertices, one gets

$$\begin{aligned} \|\Pi^\perp x\|_L^2 &\leq \sum_{r=1}^n \left(\sum_{e_{ij} \in \mathcal{E}(r)} w_{ij} (y_r(i) - y_r(j))^2 + 2 \sum_{e_{ij} \in \partial \mathcal{E}(r)} w_{ij} (y(i) - 0)^2 \right) \\ &= \sum_{r=1}^n \left(\sum_{e_{ij} \in \mathcal{E}(r)} w_{ij} (y_r(i) - y_r(j))^2 + \sum_{e_{ij} \in \partial \mathcal{E}(r)} (2w_{ij}) (y_r(i) - y_r(j))^2 \right) \\ &= \sum_{r=1}^n \|y_r\|_{L_{\mathcal{V}(r)}}^2 = \sum_{\mathcal{C} \in \mathcal{P}} \|y_r\|_{L_{\mathcal{C}}}^2. \end{aligned}$$

The second step above used the fact that $[\Pi^\perp](i) = 0$ for all $v_i \notin \mathcal{C}$. A decoupled bound can then be obtained as follows:

$$\|\Pi^\perp A\|_L^2 = \max_{a \in \mathbb{R}^{k-1}} \frac{\|S \Pi^\perp A a\|_2^2}{\|a\|_2^2} \leq \sum_{\mathcal{C} \in \mathcal{P}} \max_{a \in \mathbb{R}^{k-1}} \frac{\|\Pi_{\mathcal{C}}^\perp A a\|_{L_{\mathcal{C}}}^2}{\|a\|_2^2} = \sum_{\mathcal{C} \in \mathcal{P}} \|\Pi_{\mathcal{C}}^\perp A\|_{L_{\mathcal{C}}}^2$$

The final inequality is derived by taking the square-root of the last equation.

Appendix B. Complexity Analysis

The computational complexity of Algorithm 1 depends on the number of nodes N and edges M of G , the number of levels c , the subspace size k , as well as on how the families of candidate sets are formed. To derive worst-case bounds, I denote by $\Phi_\ell = \sum_{\mathcal{C} \in \mathcal{F}_\ell} |\mathcal{C}|$ the number of vertices in all candidate sets and by $\delta = \max_{\ell, \mathcal{C} \in \mathcal{F}_\ell} |\mathcal{C}|$ the cardinality of the maximum candidate set over all levels. Furthermore, I suppose that the per-level reduction ratio r_ℓ is a constant.

I start with some basic observations:

- *Computing A_0, \dots, A_{c-1} is possible in $\tilde{O}(ckM + k^2N + ck^3)$ operations when $V = U_k$. Each $A_{\ell-1}$ is computed once for each level. For $\ell = 1$, one needs to approximate the first k eigenpairs of L , which can be achieved in $\tilde{O}(kM)$ operations using inverse iteration as described by Vishnoi et al. (2013). For consecutive levels, forming matrix*

$B_{\ell-1}^\top L_{\ell-1} B_{\ell-1}$ takes $O(M_{\ell-1}k + N_{\ell-1}k^2)$ operations, whereas computing the pseudo-square-root $(B_{\ell-1}^\top L_{\ell-1} B_{\ell-1})^{+1/2}$ is possible in $O(k^3)$ operations. Summed up, the costs for all levels amount to $O(k \sum_{\ell=1}^c M_{\ell-1} + k^2 \sum_{\ell=2}^c N_{\ell-1} + ck^3) = O(ckM + k^2N + ck^3)$, where I used the observation that $\sum_{\ell=2}^c N_{\ell-1} = O(N)$.

- *At each level, the cost function is evaluated at most Φ_ℓ times.* One starts by computing the cost of each candidate set in \mathcal{F}_ℓ . Moreover, every \mathcal{C} added to \mathcal{P}_ℓ causes the pruning of at most $\sum_{v_i \in \mathcal{C}} (\phi_i - 1)$ other sets, where ϕ_i is the number of candidate sets that include v_i . Since \mathcal{P}_ℓ is a partitioning of $\mathcal{V}_{\ell-1}$, at most $\sum_{\mathcal{C} \in \mathcal{P}_\ell} \sum_{v_i \in \mathcal{C}} (\phi_i - 1) \leq \sum_{v_i \in \mathcal{V}_{\ell-1}} \phi_i - |\mathcal{F}_\ell| = \Phi_\ell - |\mathcal{F}_\ell|$ cost re-evaluation are needed.
- *Given $A_{\ell-1}$, each call of $\text{cost}_\ell(\mathcal{C})$ requires $O(\min\{k^2\delta + k\delta^2, k\delta^2 + \delta^3\})$ operations.* The involved matrices themselves can be easily formed since, excluding all-zero rows and columns, both $L_{\mathcal{C}}$ and $\Pi_{\mathcal{C}}^\perp$ are $|\mathcal{C}| \times |\mathcal{C}|$ matrices and one can safely restrict $A_{\ell-1}$ to be of size $|\mathcal{C}| \times k$ by deleting all rows that would have been multiplied by zero. Now, by definition, the incidence matrix $S_{\mathcal{C}}$ of $L_{\mathcal{C}}$ has at most δ columns and 2δ rows (since one can bundle all boundary weights of a vertex in \mathcal{C} in a single row). Depending on the relative size of k and δ the computation can be performed in two ways:
 - Either one forms the $k \times k$ matrix $A_{\ell-1}^\top \Pi_{\mathcal{C}}^\perp L_{\mathcal{C}} \Pi_{\mathcal{C}}^\perp A_{\ell-1}$ and approximate its spectral norm paying a total of $O(k^2\delta + k\delta^2)$.
 - Otherwise, the $2d \times 2d$ matrix $S_{\mathcal{C}} \Pi_{\mathcal{C}}^\perp A_{\ell-1} A_{\ell-1}^\top \Pi_{\mathcal{C}}^\perp S_{\mathcal{C}}^\top$ is formed and its norm is computed at a combined cost of $O(\delta^2k + \delta^3)$.
- *Maintaining \mathcal{F}_ℓ sorted incurs $O(\Phi_\ell \log |\mathcal{F}_\ell|)$ cost.* Sorting \mathcal{F}_ℓ during initialization entails $O(|\mathcal{F}_\ell| \log |\mathcal{F}_\ell|)$ operations. Inserting each \mathcal{C}' into \mathcal{F}_ℓ (see step 12) can be done in $O(\log |\mathcal{F}_\ell|)$ and, moreover, by the same argument used to bound the number of cost evaluations, at most $\Phi_\ell - |\mathcal{F}_\ell|$ such insertions can happen.
- *Other operations carry negligible cost.* In particular, by implementing marked as a binary array, checking if a vertex is marked or not can be done in constant time.

Overall, using Algorithm 2 and for $\mathbf{R} = \mathbf{U}_k$ one can coarsen a graph in $\tilde{O}(ckM + k^2N + ck^3 + \sum_{\ell=1}^c \Phi_\ell(\min\{k^2\delta + k\delta^2, k\delta^2 + \delta^3\} + \log |\mathcal{F}_\ell|))$ time, where the asymptotic notation hides poly-logarithmic factors.

References

- Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- Noga Alon and Vitali D Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38(1):73–88, 1985.
- Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W Pellegrini, Ralf S Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*, 2018.

- Joshua Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. Spectral sparsification of graphs: Theory and algorithms. *Commun. ACM*, 56(8):87–94, 2013. ISSN 0001-0782.
- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Christos Boutsidis, Prabhajan Kambadur, and Alex Gittens. Spectral clustering via the power method-provably. In *International Conference on Machine Learning*, pages 40–48, 2015.
- M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, July 2017. ISSN 1053-5888. doi: 10.1109/MSP.2017.2693418.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann Lecun. Spectral networks and locally connected networks on graphs. In *International Conference on Learning Representations (ICLR2014), CBLS, April 2014*, 2014.
- Daniele Calandriello, Ioannis Koutis, Alessandro Lazaric, and Michal Valko. Improved large-scale graph learning through ridge spectral sparsification. In *International Conference on Machine Learning*, 2018.
- Guantao Chen, George Davis, Frank Hall, Zhongshan Li, Kinnari Patel, and Michael Stewart. An interlacing result on normalized Laplacians. *SIAM Journal on Discrete Mathematics*, 18(2):353–361, 2004.
- Jie Chen and Ilya Safro. Algebraic distance on graphs. *SIAM Journal on Scientific Computing*, 33(6):3468–3490, 2011.
- Fan RK Chung. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.
- Charles Colley, Junyuan Lin, Xiaozhe Hu, and Shuchin Aeron. Algebraic multigrid for least squares problems on graphs with applications to hodgerank. In *Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2017 IEEE International*, pages 627–636. IEEE, 2017.
- Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852, 2016.
- Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. Weighted graph cuts without eigenvectors a multilevel approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11), 2007.
- Florian Dörfler and Francesco Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems*, 60(1):150–163, 2013.

- Yuguang Fang, Kenneth A Loparo, and Xiangbo Feng. Inequalities for the trace of matrix product. *IEEE Transactions on Automatic Control*, 39(12):2489–2490, 1994.
- Zvi Galil. Efficient algorithms for finding maximum matching in graphs. *ACM Computing Surveys (CSUR)*, 18(1):23–38, 1986.
- Shivam Gandhi. Improvement of the cascadic multigrid algorithm with a Gauss-Seidel smoother to efficiently compute the Fiedler vector of a graph Laplacian. *arXiv preprint arXiv:1602.04386*, 2016.
- Matan Gavish, Boaz Nadler, and Ronald R Coifman. Multiscale wavelets on trees, graphs and high dimensional data: Theory and applications to semi supervised learning. In *ICML*, pages 367–374, 2010.
- David Gleich. MatlabBGL. a matlab graph library. *Institute for Computational and Mathematical Engineering, Stanford University*, 2008.
- Francesco Grassi, Andreas Loukas, Nathanael Perraudin, and Benjamin Ricaud. A time-vertex signal processing framework: Scalable processing and meaningful representations for time-series on graphs. *IEEE Transactions on Signal Processing*, 66(3):817–829, 2017.
- Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- Bruce Hendrickson and Robert W Leland. A multi-level algorithm for partitioning graphs. *SC*, 95(28):1–14, 1995.
- Anil N Hirani, Kaushik Kalyanaraman, and Seth Watts. Graph Laplacians and least squares on graphs. In *Parallel and Distributed Processing Symposium Workshop (IPDPSW), 2015 IEEE International*, pages 812–821. IEEE, 2015.
- Yifan Hu. Efficient, high-quality force-directed graph drawing. *Mathematica Journal*, 10(1):37–71, 2005.
- Elvin Isufi, Andreas Loukas, Nathanael Perraudin, and Geert Leus. Forecasting time series with varma recursions on graphs. *arXiv preprint arXiv:1810.08581*, 2018.
- Hawoong Jeong, Sean P Mason, A-L Barabási, and Zoltan N Oltvai. Lethality and centrality in protein networks. *Nature*, 411(6833):41, 2001.
- David R Karger. Random sampling in cut, flow, and network design problems. *Mathematics of Operations Research*, 24(2):383–413, 1999.
- George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.
- David Harel Yehuda Koren. A fast multi-scale method for drawing large graphs. *Journal of graph algorithms and applications*, 6(3):179–202, 2002.

- I. Koutis, G. L. Miller, and R. Peng. Approaching optimality for solving sdd linear systems. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 235–244, Oct 2010.
- Ioannis Koutis, Gary L Miller, and David Tolliver. Combinatorial preconditioners and multilevel solvers for problems in computer vision and image processing. *Computer Vision and Image Understanding*, 115(12):1638–1646, 2011.
- Dan Kushnir, Meirav Galun, and Achi Brandt. Fast multiscale clustering and manifold identification. *Pattern Recognition*, 39(10):1876–1891, 2006.
- Stephane Lafon and Ann B Lee. Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning, and data set parameterization. *IEEE transactions on pattern analysis and machine intelligence*, 28(9):1393–1403, 2006.
- James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):37, 2014.
- Ren-Cang Li. Relative perturbation theory:(ii) eigenspace variations. Technical report, 1994.
- Jiongqian Liang, Saket Gururkar, and Srinivasan Parthasarathy. Mile: A multi-level framework for scalable graph embedding. *arXiv preprint arXiv:1802.09612*, 2018.
- Oren E Livne and Achi Brandt. Lean algebraic multigrid (LAMG): Fast graph Laplacian linear solver. *SIAM Journal on Scientific Computing*, 34(4):B499–B522, 2012.
- Andreas Loukas and Nathanaël Perraudin. Stationary time-vertex signal processing. *arXiv preprint arXiv:1611.00255*, 2016.
- Andreas Loukas and Pierre Vandergheynst. Spectrally approximating large graphs with smaller graphs. In *International Conference on Machine Learning (ICML)*, 2018.
- Lionel Martin, Andreas Loukas, and Pierre Vandergheynst. Fast approximate spectral clustering for dynamic networks. In *International Conference on Machine Learning (ICML)*, 2018.
- Ankur Moitra. *Vertex sparsification and universal rounding algorithms*. PhD thesis, Massachusetts Institute of Technology, 2011.
- Ami Paz and Gregory Schwartzman. A $(2+\epsilon)$ -approximation for maximum weight matching in the semi-streaming model. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2153–2161. SIAM, 2017.
- David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989.
- Robert Preis and Ralf Diekmann. Party-a software library for graph partitioning. *Advances in Computational Mechanics with Parallel and Distributed Processing*, pages 63–71, 1997.

- Gilles Puy, Nicolas Tremblay, Rémi Gribonval, and Pierre Vandergheynst. Random sampling of bandlimited signals on graphs. *Applied and Computational Harmonic Analysis*, 44(2):446–475, 2018.
- Dorit Ron, Ilya Safro, and Achi Brandt. Relaxation-based coarsening and multiscale graph organization. *Multiscale Modeling & Simulation*, 9(1):407–423, 2011.
- Ilya Safro, Peter Sanders, and Christian Schulz. Advanced coarsening schemes for graph partitioning. *Journal of Experimental Algorithmics (JEA)*, 19:2–2, 2015.
- Nauman Shahid, Nathanael Perraudin, Vassilis Kalofolias, Gilles Puy, and Pierre Vandergheynst. Fast robust PCA on graphs. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):740–756, 2016.
- David I Shuman, Mohammad Javad Faraji, and Pierre Vandergheynst. A multiscale pyramid transform for graph signals. *IEEE Transactions on Signal Processing*, 64(8):2119–2134, 2016.
- Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- Gilbert W Stewart. Matrix perturbation theory. 1990.
- Nicolas Tremblay and Andreas Loukas. Approximating spectral clustering via sampling: a review. *CoRR*, abs/1901.10204, 2019.
- Greg Turk and Marc Levoy. Zippered polygon meshes from range images. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 311–318. ACM, 1994.
- John C Urschel, Xiaozhe Hu, Jinchao Xu, and Ludmil T Zikatanov. A cascadic multigrid algorithm for computing the Fiedler vector of graph Laplacians. *arXiv preprint arXiv:1412.0565*, 2014.
- Nisheeth K Vishnoi et al. $Lx=b$. *Foundations and Trends® in Theoretical Computer Science*, 8(1–2):1–141, 2013.
- Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- Chris Walshaw. A multilevel algorithm for force-directed graph drawing. In *International Symposium on Graph Drawing*, pages 171–182. Springer, 2000.

- Chris Walshaw. A multilevel algorithm for force-directed graph-drawing. *Journal of Graph Algorithms and Applications*, 7(3):253–285, 2006.
- Lu Wang, Yanghua Xiao, Bin Shao, and Haixun Wang. How to partition a billion-node graph. In *Data Engineering (ICDE), 2014 IEEE 30th International Conference on*, pages 568–579. IEEE, 2014.
- Shusen Wang and Zhihua Zhang. Improving cur matrix decomposition and the Nyström approximation via adaptive sampling. *The Journal of Machine Learning Research*, 14(1): 2729–2769, 2013.
- Richard C Wilson and Ping Zhu. A study of graph spectra for comparing graphs and trees. *Pattern Recognition*, 41(9):2833–2841, 2008.
- Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2014.