Similarity-Aware Spectral Sparsification by Edge Filtering

Zhuo Feng

Department of Electrical and Computer Engineering
Michigan Technological University
Houghton, MI, 49931
zhuofeng@mtu.edu

ABSTRACT

In recent years, spectral graph sparsification techniques that can compute ultra-sparse graph proxies have been extensively studied for accelerating various numerical and graph-related applications. Prior nearly-linear-time spectral sparsification methods first extract low-stretch spanning tree from the original graph to form the backbone of the sparsifier, and then recover small portions of spectrally-critical off-tree edges to the spanning tree to significantly improve the approximation quality. However, it is not clear how many off-tree edges should be recovered for achieving a desired spectral similarity level within the sparsifier. Motivated by recent graph signal processing techniques, this paper proposes a similarity-aware spectral graph sparsification framework that leverages efficient spectral off-tree edge embedding and filtering schemes to construct spectral sparsifiers with guaranteed spectral similarity (relative condition number) level. An iterative graph densification scheme is introduced to facilitate efficient and effective filtering of off-tree edges for highly ill-conditioned problems. The proposed method has been validated using various kinds of graphs obtained from public domain sparse matrix collections relevant to VLSI CAD, finite element analysis, as well as social and data networks frequently studied in many machine learning and data mining applications.

Categories and Subject Descriptors

B.7.2 [Design Aids]: simulation—*Integrated Circuits*

General Terms

Performance, Algorithms, Verification

Keywords

Spectral graph theory, graph partitioning, iterative methods

1. INTRODUCTION

Spectral methods are playing increasingly important roles in many graph and numerical applications [22], such as scientific computing [20], numerical optimization [3], data mining [14], graph analytics [10], machine learning [7], graph signal

processing [16], and VLSI computer-aided design [9, 23]. For example, classical spectral graph partitioning (data clustering) algorithms embed original graphs into low-dimensional space using the first few nontrivial eigenvectors of graph Laplacians and subsequently perform graph partitioning (data clustering) on the low-dimensional graphs to obtain high-quality solution [14].

Recent *spectral graph sparsification* research [2, 4, 12, 14, 17, 19] allows computing nearly-linear-sized subgraphs (sparsifiers) that can robustly preserve the spectrum (i.e., eigenvalues and eigenvectors) of the original graph's Laplacian, which immediately leads to a series of theoretically nearly-linear-time numerical and graph algorithms for solving sparse matrices, graph-based semi-supervised learning (SSL), spectral graph partitioning (data clustering), and max-flow problems [3, 11, 19, 20]. For example, sparsified circuit networks allow for developing more scalable computer-aided (CAD) design algorithms for designing large VLSI systems [9, 23]; sparsified social (data) networks enable to more efficiently understand and analyze large social (data) networks [22]; sparsified matrices can be immediately leveraged to accelerate the solution computation of large linear system of equations [24]. To this end, a spectral sparsification algorithm leveraging an edge sampling scheme that sets sampling probabilities proportional to edge effective resistances (of the original graph) has been proposed in [17].

A practically-efficient, nearly-linear complexity spectral graph sparsification algorithm has been recently introduced in [9], which first extracts a “spectrally critical” spanning tree subgraph as a backbone of the sparsifier, and subsequently recovers a small portion of dissimilar “spectrally critical” off-tree edges to the spanning tree. However, in many scientific computing and graph-related applications, it is important to compute spectral graph sparsifiers of desired spectral similarity level: introducing too few edges may lead to poor approximation of the original graph, whereas too many edges can result in high computational complexity. For example, when using a preconditioned conjugate gradient (PCG) solver to solve a symmetric diagonally dominant (SDD) matrix for multiple right-hand-side (RHS) vectors, it is hoped the PCG solver would converge to a good solution as quickly as possible, which usually requires the sparsifier (preconditioner) to be highly spectrally-similar to the original problem; on the other hand, in many graph partitioning tasks, only the Fiedler vector (the first nontrivial eigenvector) of graph Laplacian is needed [18], so even a sparsifier with much lower spectral similarity will suffice.

This paper introduces a similarity-aware spectral graph sparsification framework that leverages efficient spectral off-tree edge embedding and filtering schemes to construct spectral sparsifiers with guaranteed spectral similarity. The contribution of this work has been summarized as follows:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC'18, June 24–29, 2018, San Francisco, CA, USA
Copyright 2018 ACM 978-1-4503-5700-5/18/06 ...\$15.00
<https://doi.org/10.1145/3195970.3196114>.

1. We present a similarity-aware spectral graph sparsification framework by leveraging spectral off-tree edge embedding and filtering schemes that have been motivated by recent graph signal processing techniques [16].
2. An iterative graph densification procedure is proposed to incrementally improve the approximation of the sparsifier, which enables to flexibly trade off the complexity and spectral similarity of the sparsified graph.
3. Extensive experiments have been conducted to validate the proposed method in various numerical and graph-related applications, such as solving sparse SDD matrices, and spectral graph partitioning, as well as simplification of large social and data networks.

2. SPECTRAL GRAPH SPARSIFICATION

Consider a graph $G = (V, E, w)$ with V denoting the vertex (data point) set of the graph, E denoting the edge set of the graph, and w denoting a weight (similarity) function that assigns positive weights to all edges. The graph Laplacian $\mathbf{L}_G$ of G is an SDD matrix defined as follows:

$$\mathbf{L}_G(p, q) = \begin{cases} -w(p, q) & \text{if } (p, q) \in E \\ \sum_{(p, t) \in E} w(p, t) & \text{if } (p, q) = (p, p) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Spectral graph sparsification [19] aims to preserve the original graph spectrum within ultra-sparse subgraphs (graph sparsifiers), which allows preserving not only cuts in the graph but also eigenvalues and eigenvectors of the original graph Laplacian, distances (e.g. effective resistances) between vertices, low-dimensional graph embedding, etc. Two graphs G and P are said to be σ -*spectrally similar* if for all real vectors $\mathbf{x} \in \mathbb{R}^V$ their quadratic forms satisfy:

$$\frac{\mathbf{x}^T \mathbf{L}_P \mathbf{x}}{\sigma} \leq \mathbf{x}^T \mathbf{L}_G \mathbf{x} \leq \sigma \mathbf{x}^T \mathbf{L}_P \mathbf{x}. \quad (2)$$

Define the relative condition number to be $\kappa(\mathbf{L}_G, \mathbf{L}_P) = \lambda_{\max}/\lambda_{\min}$, where $\lambda_{\max}$ and $\lambda_{\min}$ denote the largest and smallest generalized eigenvalues satisfying:

$$\mathbf{L}_G \mathbf{u} = \lambda \mathbf{L}_P \mathbf{u}, \quad (3)$$

with $\mathbf{u}$ denoting the generalized eigenvector of λ . It can be further shown that $\kappa(\mathbf{L}_G, \mathbf{L}_P) \leq \sigma^2$, which indicates that a smaller relative condition number or σ^2 corresponds to a higher spectral similarity. Obviously, we can simply use σ^2 to denote the upper bound of the relative condition number.

3. SIMILARITY-AWARE SPECTRAL SPARSIFICATION BY EDGE FILTERING

3.1 Overview of Our Approach

The overview of the proposed method for similarity-aware spectral sparsification of undirected graphs has been summarized as follows. For a given input graph, the following key procedures are involved in the proposed algorithm flow: (a) low-stretch spanning tree [1, 8] extraction based on its original graph Laplacian; (b) spectral (generalized eigenvalue) embedding and filtering of off-tree edges by leveraging the recent spectral perturbation analysis framework [9]; (c) incremental sparsifier improvement (graph densification) by gradually adding small portions of dissimilar off-tree edges to the spanning tree. Fig. 1 shows the spectral drawings [10] of an airfoil graph [6] as well as its spectrally-similar subgraph computed by the proposed similarity-aware spectral sparsification algorithm.

Figure 1: Two spectrally-similar airfoil graphs.

In the rest of this paper, we assume that $G = (V, E, w)$ is a weighted, undirected and connected graph, whereas $P = (V, E_s, w_s)$ is its sparsifier. To simplify the our analysis, we assume the edge weights in the sparsifier remain the same as the original ones, though edge re-scaling schemes [19] can be applied to further improve the approximation. The descending eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ are denoted by $\lambda_{\max} = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 1$, where $\mathbf{L}_P^+$ denotes the Moore-Penrose pseudoinverse of $\mathbf{L}_P$.

3.2 Spectral Embedding of Off-Tree Edges

It has been shown that there are not too many large generalized eigenvalues for spanning-tree sparsifiers [21]: $\mathbf{L}_P^+ \mathbf{L}_G$ has at most k generalized eigenvalues greater than $\frac{\text{st}_P(G)}{k}$, where $\text{st}_P(G)$ denotes the total stretch of the spanning-tree subgraph P with respect the original graph G . Recent research results show that every undirected graph has a low-stretch spanning tree (LSST) such that [1, 8]:

$$O(m \log n \log \log n) \geq \text{st}_P(G) = \text{Trace}(\mathbf{L}_P^+ \mathbf{L}_G) = \sum_{i=1}^n \lambda_i, \quad (4)$$

where $m = |E|$ and $n = |V|$. As a result, it is possible to construct an ultra-sparse yet spectrally similar sparsifier by recovering only a small portion of important off-tree edges to the spanning tree: for example, σ -similar spectral sparsifiers with $O(\frac{m \log n \log \log n}{\sigma^2})$ off-tree edges can be computed efficiently using perturbation-based method [9].

To identify important off-tree edges the following generalized eigenvalue perturbation analysis is considered [9]:

$$\mathbf{L}_G(\mathbf{u}_i + \delta \mathbf{u}_i) = (\lambda_i + \delta \lambda_i)(\mathbf{L}_P + \delta \mathbf{L}_P)(\mathbf{u}_i + \delta \mathbf{u}_i), \quad (5)$$

where a perturbation matrix $\delta \mathbf{L}_P$ is applied for the inclusion of extra off-tree edges into $\mathbf{L}_P$ and results in perturbed generalized eigenvalues and eigenvectors $\lambda_i + \delta \lambda_i$ and $\mathbf{u}_i + \delta \mathbf{u}_i$ for $i = 1, \dots, n$, respectively. The key to effective spectral sparsification is to identify the key off-tree edges that will result in the greatest reduction in dominant generalized eigenvalues. To this end, the following scheme for embedding generalized eigenvalues into each off-tree edge is adopted in this work [9]:

Step 1: Start with an initial random vector $\mathbf{h}_0 = \sum_{i=1}^n \alpha_i \mathbf{u}_i$,

where $\mathbf{u}_i$ are the $\mathbf{L}_P$ -orthogonal generalized eigenvectors of $\mathbf{L}_P^+ \mathbf{L}_G$ that satisfy $\mathbf{u}_i^T \mathbf{L}_P \mathbf{u}_j = 1$ for $i = j$, and 0 for $i \neq j$;

Step 2: Perform t -step generalized power iterations with $\mathbf{h}_0$ to obtain $\mathbf{h}_t = (\mathbf{L}_P^+ \mathbf{L}_G)^t \mathbf{h}_0 = \sum_{i=1}^n \alpha_i \lambda_i^t \mathbf{u}_i$; $\mathbf{h}_t$ will be a good approximation of dominant eigenvectors;

Step 3: Compute the Laplacian quadratic form for $\delta \mathbf{L}_{P, \max} =$

$\mathbf{L}_G - \mathbf{L}_P$ with $\mathbf{h}_t$:

$$\begin{aligned} Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_t) &= \mathbf{h}_t^T \delta \mathbf{L}_P, \max \mathbf{h}_t = \sum_{i=1}^n (\alpha_i \lambda_i^t)^2 (\lambda_i - 1) \\ &= \sum_{(p,q) \in E \setminus E_s} w_{p,q} \sum_{i=1}^n \alpha_i^2 \lambda_i^{2t} (\mathbf{u}_i^T \mathbf{e}_{p,q})^2 = \sum_{(p,q) \in E \setminus E_s} \text{heat}_{(p,q)}, \end{aligned} \quad (6)$$

where $\delta \mathbf{L}_P, \max$ denotes the perturbation of $\mathbf{L}_P$ including all off-tree edges, $\mathbf{e}_{p,q} \in \mathbb{R}^n$ denotes the vector that has the p -th element being 1, the q -th element being -1 and others being 0, and $\text{heat}_{(p,q)}$ denotes the edge Joule heat of the off-tree edge (p,q) . The amplitude of $Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_t)$ reflects the spectral similarity between graphs P and G : larger $Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_t)$ indicates greater σ^2 and thus lower spectral similarity. More importantly, (6) allows embedding generalized eigenvalues into the Laplacian quadratic form of each off-tree edge and ranking each off-tree edge according to its edge Joule heat (spectral criticality): recovering the off-tree edges with largest $\text{heat}_{(p,q)}$ will most significantly decrease the largest generalized eigenvalues. In practice, using a small number (e.g. $t = 2$) of generalized power iterations will suffice for spectral edge embedding purpose.

3.3 “Spectrally-Unique” Off-Tree Edges

To simplify the following analysis, we define a “spectrally unique” off-tree edge e_i to be the one that connects to vertices p_i and q_i , and only impacts a single large generalized eigenvalue λ_i . Then the truncated version of (6) including the top k dominant “spectrally-unique” off-tree edges for fixing the top k largest eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ can be expressed as follows for $\lambda_i \gg 1$:

$$Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_t) \approx \sum_{i=1}^k w_{p_i, q_i} \alpha_i^2 \lambda_i^{2t} (\mathbf{u}_i^T \mathbf{e}_{p_i, q_i})^2 \approx \sum_{i=1}^k \alpha_i^2 \lambda_i^{2t+1}. \quad (7)$$

Since each off-tree edge only impacts one generalized eigenvalue, we can express $\mathbf{e}_{p_i, q_i} = \gamma_i \mathbf{L}_P \mathbf{u}_i$ according to (6), which leads to:

$$\mathbf{u}_j^T \mathbf{e}_{p_i, q_i} = \begin{cases} \gamma_i, & i = j, \\ 0, & i \neq j. \end{cases} \quad (8)$$

Then the effective resistance of edge e_i in P becomes:

$$R_{e_i}^{\text{eff}} = \mathbf{e}_{p_i, q_i}^T \mathbf{L}_P^+ \mathbf{e}_{p_i, q_i} = \gamma_i^2 \mathbf{u}_i^T \mathbf{L}_P \mathbf{u}_i = \gamma_i^2, \quad (9)$$

which immediately leads to:

$$Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_t) \approx \sum_{i=1}^k \alpha_i^2 \lambda_i^{2t} w_{p_i, q_i} R_{e_i}^{\text{eff}} \approx \sum_{i=1}^k \alpha_i^2 \lambda_i^{2t+1}. \quad (10)$$

Since the stretch of off-tree edge e_i is computed by $\text{st}_P(e_i) = w_{p_i, q_i} R_{e_i}^{\text{eff}}$, (10) also indicates that $\text{st}_P(e_i) \approx \lambda_i$ holds for “spectrally-unique” off-tree edges. Consequently, the key off-tree edges identified by (6) or (10) will have the largest stretch values and therefore most significantly impact the largest eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$. (10) also can be considered as a randomized version of $\text{Trace}(\mathbf{L}_P^+ \mathbf{L}_G)$ that is further scaled up by a factor of λ_i^{2t} .

3.4 Spectral Sparsification as A Graph Filter

Although (6) and (10) provide a spectral ranking for each off-tree edge, it is not clear how many off-tree edges should be recovered to the spanning tree for achieving a desired spectral similarity level. To this end, we introduce a simple yet effective spectral off-tree edge filtering scheme motivated by recent graph signal processing techniques [16]. To more efficiently analyze signals on general undirected

graphs, graph signal processing techniques have been extensively studied recently [16]. There is a clear analogy between traditional signal processing based on classical Fourier analysis and graph signal processing: 1) the signals at different time points in classical Fourier analysis correspond to the signals at different nodes in an undirected graph; 2) the more slowly oscillating functions in time domain correspond to the graph Laplacian eigenvectors associated with lower eigenvalues and more slowly varying (smoother) components across the graph. A comprehensive review of fundamental signal processing operations, such as filtering, translation, modulation, dilation, and down-sampling to the graph setting has been provided in [16].

Spectral sparsification aims to maintain a simplest sub-graph sufficient for preserving the slowly-varying or “low-frequency” signals on graphs, which therefore can be regarded as a “low-pass” graph filter. In other words, such spectrally sparsified graphs will be able to preserve the eigenvectors associated with low eigenvalues more accurately than high eigenvalues, and thus will retain “low-frequency” graph signals sufficiently well, but not so well for highly-oscillating (signal) components due to the missing edges.

In practice, preserving the spectral (structural) properties of the original graph within the spectral sparsifier is key to design of many fast numerical and graph-related algorithms [3, 11, 17, 20]. For example, when using spectral sparsifier as a preconditioner in preconditioned conjugate gradient (PCG) iterations, the convergence rate only depends on the spectral similarity (or relative condition number) for achieving a desired accuracy level, while in spectral graph partitioning and data clustering tasks only the first few eigenvectors associated with the smallest nontrivial eigenvalues of graph Laplacian are needed [14, 18].

3.5 Off-Tree Edge Filtering with Joule Heat

To only recover the off-tree edges that are most critical for achieving the desired spectral similarity level, we propose the following scheme for truncating “spectrally-unique” off-tree edges based on each edge’s Joule heat. For a spanning-tree preconditioner, since there will be at most k generalized eigenvalues that are greater than $\text{st}_P(G)/k$, the following simple yet nearly worst-case generalized eigenvalue distribution can be assumed:

$$\lambda_i = \frac{2\lambda_{\max}}{i+1} = \frac{\text{st}_P(G)}{i+1}, i \geq 1. \quad (11)$$

To most economically select the top- k “spectrally-unique” off-tree edges that will dominantly impact the top- k largest generalized eigenvalues, the following sum of quadratic forms (Joule heat levels) can be computed based on (10) by performing t -step generalized power iterations with r multiple random vectors $\mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,r}$:

$$Q_{\delta \mathbf{L}_P, \max}(\mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,r}) \approx \sum_{j=1}^r \sum_{i=1}^k (\alpha_{i,j})^2 \left(\frac{2\lambda_{\max}}{i+1} \right)^{2t+1}. \quad (12)$$

The goal is to select top k “spectrally-unique” off-tree edges for fixing the top k largest generalized eigenvalues such that the resulting upper bound of the relative condition number will become $\sigma^2 = \frac{\tilde{\lambda}_{\max}}{\tilde{\lambda}_{\min}}$, where $\tilde{\lambda}_{\max}$ and $\tilde{\lambda}_{\min}$ denote the largest and smallest eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ after adding top- k “spectrally-unique” off-tree edges. Then we have:

$$k = 2\lambda_{\max}/\tilde{\lambda}_{\max} - 1. \quad (13)$$

When using multiple random vectors for computing (12), it

is expected that $\sum_{j=1}^r \alpha_{k,j}^2 \approx \sum_{j=1}^r \alpha_{1,j}^2$, which allows us to define the normalized edge Joule heat θ_k for the k -th “spectrally-unique” off-tree edge through the following simplifications:

$$\theta_k = \frac{\text{heat}_{\lambda_k}}{\text{heat}_{\lambda_1}} = \left(\frac{\sum_{j=1}^r \alpha_{k,j}^2}{\sum_{j=1}^r \alpha_{1,j}^2} \right) \left(\frac{\tilde{\lambda}_{\max}}{\lambda_{\max}} \right)^{2t+1} \approx \left(\frac{\sigma^2 \tilde{\lambda}_{\min}}{\lambda_{\max}} \right)^{2t+1}. \quad (14)$$

The key idea of the proposed similarity-aware spectral sparsification is to leverage the normalized Joule heat (14) as a threshold for filtering off-tree edges: only the off-tree edges with normalized Joule heat values greater than θ_k will be selected for inclusion into the spanning tree for achieving the desired spectral similarity (σ) level. Although the above scheme is derived for filtering “spectrally-unique” off-tree edges, general off-tree edges also can be filtered using similar strategies. Since adding the off-tree edges with largest Joule heat to the subgraph will mainly impact the largest generalized eigenvalues but not the smallest ones, we will assume $\tilde{\lambda}_{\min} \approx \lambda_{\min}$, and use the following edge truncation scheme for filtering general off-tree edges: the off-tree edge (p, q) will be included into the sparsifier if its normalized Joule heat value is greater than the threshold determined by:

$$\theta_{(p,q)} = \frac{\text{heat}_{(p,q)}}{\text{heat}_{\max}} \geq \theta_\sigma \approx \left(\frac{\sigma^2 \lambda_{\min}}{\lambda_{\max}} \right)^{2t+1}, \quad (15)$$

where θ_σ denotes the threshold for achieving the σ -spectral similarity in the sparsifier, and $\text{heat}_{\max}$ denotes the maximum Joule heat of all off-tree edges computed by (6) with multiple initial random vectors.

3.6 Estimation of Extreme Eigenvalues

To achieve the above spectral off-tree edge filtering scheme, we need to compute θ_σ in (15) that further requires to estimate the extreme eigenvalues $\lambda_{\max}$ and $\lambda_{\min}$ of $\mathbf{L}_P^+ \mathbf{L}_G$. In this work, we propose the following efficient methods for computing these extreme generalized eigenvalues.

3.6.1 Estimating $\lambda_{\max}$ via Power Iterations

Since generalized power iterations converge at a geometric rate determined by the separation of the two largest generalized eigenvalues $\lambda_{\max} = \lambda_1 > \lambda_2$, the error of the estimated eigenvalue will decrease quickly when $|\lambda_2/\lambda_1|$ is small. It has been shown that the largest eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ are well separated from each other [21], which thus leads to very fast convergence of generalized power iterations for estimating λ_1 . To achieve scalable performance of power iterations, we can adopt recent graph-theoretic algebraic multigrid (AMG) methods for solving the sparsified Laplacian matrix $\mathbf{L}_P$ [13, 24].

3.6.2 Estimating $\lambda_{\min}$ via Node Coloring

Since the smallest eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ are crowded together [21], using (shifted) inverse power iterations may not be efficient due to the extremely slow convergence rate. To the extent of our knowledge, none of existing eigenvalue decomposition methods can efficiently compute $\lambda_{\min}$.

This work exploits the following Courant-Fischer theorem for generalized eigenvalue problems:

$$\lambda_{\min} = \min_{|\mathbf{x}| \neq 0} \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{L}_P \mathbf{x}}, \quad (16)$$

where $\mathbf{x}$ is also required to be orthogonal to the all-one vector. (16) indicates that if we can find a vector $\mathbf{x}$ that minimizes the ratio between the quadratic forms of the original

and sparsified Laplacians, $\lambda_{\min}$ can be subsequently computed. By restricting the values in $\mathbf{x}$ to be only 1 or 0, which can be considered as assigning one of the two colors to each node in graphs G and P , the following simplifications can be made:

$$\lambda_{\min} \approx \min_{\substack{|\mathbf{x}| \neq 0 \\ x(i) \in \{0,1\}}} \frac{\mathbf{x}^T \mathbf{L}_G \mathbf{x}}{\mathbf{x}^T \mathbf{L}_P \mathbf{x}} = \min_{\substack{|\mathbf{x}| \neq 0 \\ x(i) \in \{0,1\}}} \frac{\sum_{x(p) \neq x(q), (p,q) \in E} w_{pq}}{\sum_{x(p) \neq x(q), (p,q) \in E_s} w_{pq}}, \quad (17)$$

which will always allow estimating an upper bound for $\lambda_{\min}$. To this end, we first initialize all nodes with 0 value and subsequently try to find a node p such that the ratio between quadratic forms can be minimized:

$$\lambda_{\min} \approx \min_{p \in V} \frac{L_G(p, p)}{L_P(p, p)}. \quad (18)$$

The above procedure for estimating $\lambda_{\min}$ only requires finding the node with the smallest node degree ratio and thus can be easily implemented and efficiently performed for even very graph problems. Our results for real-world graphs show that the proposed method is highly efficient and can reasonably estimate the smallest generalized eigenvalues when compared with existing generalized eigenvalue methods [15].

3.7 Iterative Graph Densification

To achieve more effective edge filtering for similarity-aware spectral graph sparsification, we propose to iteratively recover off-tree edges to the sparsifier through an incremental graph densification procedure. Each densification iteration adds a small portion of “filtered” off-tree edges to the latest spectral sparsifier, while the spectral similarity is estimated to determine if more off-tree edges are needed. The i -th graph densification iteration includes the following steps:

1. Update the subgraph Laplacian matrix $\mathbf{L}_P$ as well as its solver by leveraging recent graph-theoretic algebraic multigrid methods [13, 24];
2. Estimate the spectral similarity by computing $\lambda_{\max}$ and $\lambda_{\min}$ using the methods described in Section 3.6;
3. If the spectral similarity is not satisfactory, continue with the following steps; otherwise, terminate the subgraph densification procedure.
4. Perform t -step generalized power iterations with $O(\log |V|)$ random vectors to compute the sum of Laplacian quadratic forms (12);
5. Rank and filter each off-tree edge according to its normalized Joule heat value using the threshold θ_σ in (15);
6. Check the similarity of each selected off-tree edge and only add dissimilar edges to the latest sparsifier.

4. EXPERIMENTAL RESULTS

The proposed spectral sparsification algorithm has been implemented in C++¹. The test cases used in this paper have been selected from a great variety of matrices that have been used in circuit simulation, finite element analysis, machine learning and data mining applications. If the original matrix is not a graph Laplacian, it will be converted into a graph Laplacian by setting each edge weight using the absolute value of each nonzero entry in the lower triangular matrix; if edge weights are not available in the original matrix file, a unit edge weight will be assigned to all edges.

¹<https://sites.google.com/mtu.edu/zhufeng-graphspar>

Table 1: Results of extreme eigenvalue estimations.

Test Cases	λ_{min}	$\tilde{\lambda}_{min}$	$\epsilon_{\lambda_{min}}$	λ_{max}	$\tilde{\lambda}_{max}$	$\epsilon_{\lambda_{max}}$
fe_rotor	1.34	1.40	4.4%	120.9	116.7	3.5%
pdb1HYS	1.71	1.89	10.5%	120.6	113.2	6.1%
bcsstk36	1.18	1.27	7.6%	96.0	92.4	3.8%
brack2	1.15	1.20	4.3%	92.6	90.3	2.5%
raefsky3	1.13	1.25	10.5%	84.4	82.7	2.0%

Figure 2: Spectral edge ranking and filtering by normalized Joule heat of off-tree edges for *G2_circuit* (left) and *Thermal1* (right) test cases [6] with top off-tree edges highlighted in red rectangles.

All of our experiments have been conducted using a single CPU core of a computing platform running 64-bit RHEW 7.2 with a 2.67GHz 12-core CPU and 50 GB memory.

4.1 Estimation of Extreme Eigenvalues

In Table 1, the extreme generalized eigenvalues ($\tilde{\lambda}_{min}$ and $\tilde{\lambda}_{max}$) estimated by the proposed methods (Section 3.6) are compared with the ones (λ_{min} and λ_{max}) computed by the “eigs” function in Matlab for sparse matrices in [6], while the relative errors ($\epsilon_{\lambda_{min}}$ and $\epsilon_{\lambda_{max}}$) are also shown. $\tilde{\lambda}_{max}$ is estimated using less than ten generalized power iterations.

We also illustrate the results of spectral edge ranking and filtering according to Joule heat levels computed by one-step generalized power iteration using (6) in Fig. 2 for two sparse matrices in [6]. The thresholds of normalized edge Joule heat values required for spectral edge filtering are labeled using red dash lines. It is observed in Fig. 2 there is a sharp change of the top normalized edge Joule heat values, which indicates that there are not many large eigenvalues of $\mathbf{L}_P^+ \mathbf{L}_G$ in both cases and agrees well with the prior theoretical analysis [21].

4.2 A Scalable Sparse SDD Matrix Solver

The spectral sparsifier obtained by the proposed similarity-aware algorithm is also leveraged as a preconditioner in a PCG solver. The RHS input vector b is generated randomly and the solver is set to converge to an accuracy level $\|Ax - b\| < 10^{-3}\|b\|$ for all test cases. “ $|V|$ ” and “ $|E|$ ” denote the numbers of nodes and edges in the original graph, whereas “ $|E_{\sigma^2}|$ ”, “ N_{σ^2} ” and “ T_{σ^2} ” denote the number of edges in the sparsifier, the number of PCG iterations required for converging to the desired accuracy level, and the total time of graph sparsification for achieving the spectral similarity of σ^2 , respectively. As observed in all test cases, there are very clear trade-offs between the graph density, computation time, and spectral similarity for all spectral sparsifiers extracted using the proposed method: sparsifiers with higher spectral similarities (smaller σ^2) allow converging to the required solution accuracy level in much fewer PCG iterations, but need to retain more edges in the subgraphs and thus require longer time to compute (sparsify).

Table 2: Results of iterative SDD matrix solver.

Graphs	$ V $	$ E $	$\frac{ E_{50} }{ V }$	N_{50}	T_{50}	$\frac{ E_{200} }{ V }$	N_{200}	T_{200}
G3_circuit	1.6E6	3.0E6	1.11	21	20s	1.05	37	8s
thermal2	1.2E6	3.7E6	1.14	20	23s	1.06	36	9s
ecology2	1.0E6	2.0E6	1.14	20	16s	1.06	40	5s
tmt_sym	0.7E6	2.2E6	1.21	19	16s	1.14	38	4s
paraboli_fem	0.5E6	1.6E6	1.22	18	16s	1.09	38	3s

Table 3: Results of spectral graph partitioning.

Test Cases	$ V $	$\frac{ V_+ }{ V_- }$	T_D (M_D)	T_I (M_I)	Rel.Err.
G3_circuit	1.6E6	1.35	52.3s (2.3G)	7.6s (0.3G)	2.2E-2
thermal2	1.2E6	1.00	13.0s (0.9G)	3.0s (0.2G)	6.8E-4
ecology2	1.0E6	1.03	12.1s (0.7G)	3.4s (0.2G)	8.9E-3
tmt_sym	0.7E6	0.99	10.2s (0.6G)	1.9s (0.1G)	2.1E-2
paraboli_fem	0.5E6	0.98	8.8s (0.4G)	2.4s (0.1G)	3.9E-2
mesh_1M	1.0E6	1.01	10.2s (0.7G)	1.7s (0.2G)	3.3E-3
mesh_4M	4.5E6	0.99	49.6s (3.0G)	8.2s (0.7G)	7.5E-3
mesh_9M	9.0E6	0.99	138.5s (6.9G)	13.3s (1.5G)	7.8E-4

4.3 A Scalable Spectral Graph Partitioner

It has been shown that by applying only a few inverse power iterations, the approximate Fiedler vector (u_f) that corresponds to the smallest nonzero eigenvalue of the (normalized) graph Laplacian matrix can be obtained for obtaining high-quality graph partitioning solution [20]. Therefore, using the spectral sparsifiers computed by the proposed spectral sparsification algorithm can immediately accelerate the PCG solver for inverse power iterations, leading to scalable performance for graph partitioning problems [20]. In fact, if the spectral sparsifier is already a good approximation of the original graph, its Fiedler vector can be directly used for partitioning the original graph.

We implement the accelerated spectral graph partitioning algorithm, and test it with sparse matrices in [6] and several 2D mesh graphs synthesized with random edge weights. As shown in Table 3, the graphs associated with sparse matrices have been partitioned into two pieces using sign cut method [18] according to the approximate Fiedler vectors computed by a few steps of inverse power iterations. The direct solver [5] and the preconditioned iterative solver are invoked within each inverse power iteration for updating the approximate Fiedler vectors u_f and $\tilde{u}_f$, respectively. “ $\frac{|V_+|}{|V_-|}$ ” denotes the ratio of nodes assigned with positive and negative signs according to the approximate Fiedler vector, and “Rel.Err.” denotes the relative error of the proposed solver compared to the direct solver computed by $\frac{|V_{dif}|}{|V|}$, where $|V_{dif}|$ denotes the number of nodes with different signs in u_f and $\tilde{u}_f$. “ T_D ” (“ T_I ”) and “ M_D ” (“ M_I ”) denote the total solution time (excluding sparsification time) and memory cost of the direct (iterative) method. We extract sparsifiers with $\sigma^2 \leq 200$ for all test cases.

4.4 Sparsification of Other Complex networks

As shown in Table 4, a few finite element, protein, data and social networks have been spectrally sparsified to achieve $\sigma^2 \approx 100$ using the proposed similarity-aware method. “ T_{tot} ”

Table 4: Results of complex network sparsification.

Test Cases	$ V $	$ E $	T_{tot}	$\frac{ E }{ E_s }$	$\frac{\lambda_1}{\lambda_{ V }}$	$T_{eig}^o(T_{eig}^s)$
fe_tooth	7.8E4	4.5E5	3.0s	5×	8,050×	14.5s (2.7s)
appu	1.4E4	9.2E5	5.4s	25×	13,624×	2,400s (15s)
coAuthorsDBLP	3.0E5	1.0E6	7.2s	3×	1,364×	2,047s (36s)
auto	4.5E5	3.3E6	29.0s	5×	48,190×	N/A (54s)
RCV-80NN	1.9E5	1.2E7	46.5s	36×	28,410×	N/A (170s)

is the total time for extracting the sparsifier, " $\frac{\lambda_1}{\lambda_1}$ " denotes the ratio of the largest generalized eigenvalues before and after adding off-tree edges into the spanning tree sparsifier, and $T_{\text{eig}}^o(T_{\text{eig}}^s)$ denotes the time for computing the first ten eigenvectors of the original (sparsified) graph Laplacians using the "eigs" function in Matlab. Since spectral sparsifiers can well approximate the spectral (structural) properties of the original graph, the sparsified graphs can be leveraged for accelerating many numerical and graph-related tasks. For example, spectral clustering (partitioning) using the original "RCV-80NN" (80-nearest-neighbor) graph can not be performed on our server with 50GB memory, while it only takes a few minutes using the sparsified one.

5. CONCLUSIONS

For the first time, this paper introduces a similarity-aware spectral graph sparsification framework that can be immediately leveraged to develop fast numerical and graph-related algorithms. Motivated by recent graph signal processing concepts, an iterative graph densification procedure based on spectral embedding and filtering of off-tree edges has been proposed for extracting ultra-sparse yet spectrally-similar graph sparsifiers, which enables to flexibly trade off the complexity and spectral similarity of the sparsified graph in numerical and graph-related applications. Extensive experimental results have confirmed the effectiveness and scalability of an iterative matrix solver and a spectral graph partitioning algorithm for a variety of large-scale, real-world graph problems, such as VLSI and finite element analysis problems, as well as data and social networks.

6. ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under Grants CCF-1350206 and CCF-1618364.

7. REFERENCES

- [1] I. Abraham and O. Neiman. Using petal-decompositions to build a low stretch spanning tree. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing (STOC)*, pages 395–406. ACM, 2012.
- [2] J. Batson, D. Spielman, and N. Srivastava. Twice-Ramanujan Sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.
- [3] P. Christiano, J. Kelner, A. Madry, D. Spielman, and S. Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proc. ACM STOC*, pages 273–282, 2011.
- [4] M. B. Cohen, J. Kelner, J. Peebles, R. Peng, A. B. Rao, A. Sidford, and A. Vladu. Almost-linear-time algorithms for markov chains and new spectral primitives for directed graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 410–419. ACM, 2017.
- [5] T. Davis. *CHOLMOD: sparse supernodal Cholesky factorization and update/downdate*. [Online]. Available: <http://www.cise.ufl.edu/research/sparse/cholmod/>, 2008.
- [6] T. Davis and Y. Hu. The university of florida sparse matrix collection. *ACM Trans. on Math. Soft. (TOMS)*, 38(1):1, 2011.
- [7] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852, 2016.
- [8] M. Elkin, Y. Emek, D. Spielman, and S. Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 38(2):608–628, 2008.
- [9] Z. Feng. Spectral graph sparsification in nearly-linear time leveraging efficient spectral perturbation analysis. In *Design Automation Conference (DAC), 2016 53rd ACM/EDAC/IEEE*, pages 1–6. IEEE, 2016.
- [10] Y. Koren. On spectral graph drawing. In *International Computing and Combinatorics Conference*, pages 496–508. Springer, 2003.
- [11] I. Koutis, G. Miller, and R. Peng. Approaching Optimality for Solving SDD Linear Systems. In *Proc. IEEE FOCS*, pages 235–244, 2010.
- [12] Y. T. Lee and H. Sun. An SDP-based Algorithm for Linear-sized Spectral Sparsification. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 678–687, New York, NY, USA, 2017. ACM.
- [13] O. Livne and A. Brandt. Lean algebraic multigrid (LAMG): Fast graph Laplacian linear solver. *SIAM Journal on Scientific Computing*, 34(4):B499–B522, 2012.
- [14] R. Peng, H. Sun, and L. Zanetti. Partitioning well-clustered graphs: Spectral clustering works. In *Proceedings of The 28th Conference on Learning Theory (COLT)*, pages 1423–1455, 2015.
- [15] Y. Saad. *Numerical Methods for Large Eigenvalue Problems: Revised Edition*, volume 66. Siam, 2011.
- [16] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- [17] D. Spielman and N. Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [18] D. Spielman and S. Teng. Spectral partitioning works: Planar graphs and finite element meshes. In *Foundations of Computer Science (FOCS), 1996. Proceedings., 37th Annual Symposium on*, pages 96–105. IEEE, 1996.
- [19] D. Spielman and S. Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- [20] D. Spielman and S. Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.
- [21] D. Spielman and J. Woo. A note on preconditioning by low-stretch spanning trees. *arXiv preprint arXiv:0903.2816*, 2009.
- [22] S.-H. Teng. Scalable algorithms for data and network analysis. *Foundations and Trends® in Theoretical Computer Science*, 12(1–2):1–274, 2016.
- [23] Z. Zhao and Z. Feng. A spectral graph sparsification approach to scalable vectorless power grid integrity verification. In *Proceedings of the 54th Annual Design Automation Conference 2017*, page 68. ACM, 2017.
- [24] Z. Zhao, Y. Wang, and Z. Feng. SAMG: Sparsified Graph Theoretic Algebraic Multigrid for Solving Large Symmetric Diagonally Dominant (SDD) Matrices. In *Proceedings of the 36th International Conference on Computer-Aided Design (ICCAD)*. ACM, 2017.