Efficient Sparsification of Simplicial Complexes via Local Densities of States

Anton Savostianov
Computational Network Science, RWTH Aachen, Germany
savostianov@cs.rwth-aachen.de

Michael T. Schaub
Computational Network Science, RWTH Aachen, Germany
schaub@cs.rwth-aachen.de

Nicola Guglielmi
Gran Sasso Science Institute
L'Aquila, Italy
nicola.guglielmi@gssi.it

Francesco Tudisco
Maxwell Institute & School of Mathematics, University of Edinburgh,
Edinburgh, UK
f.tudisco@ed.ac.uk

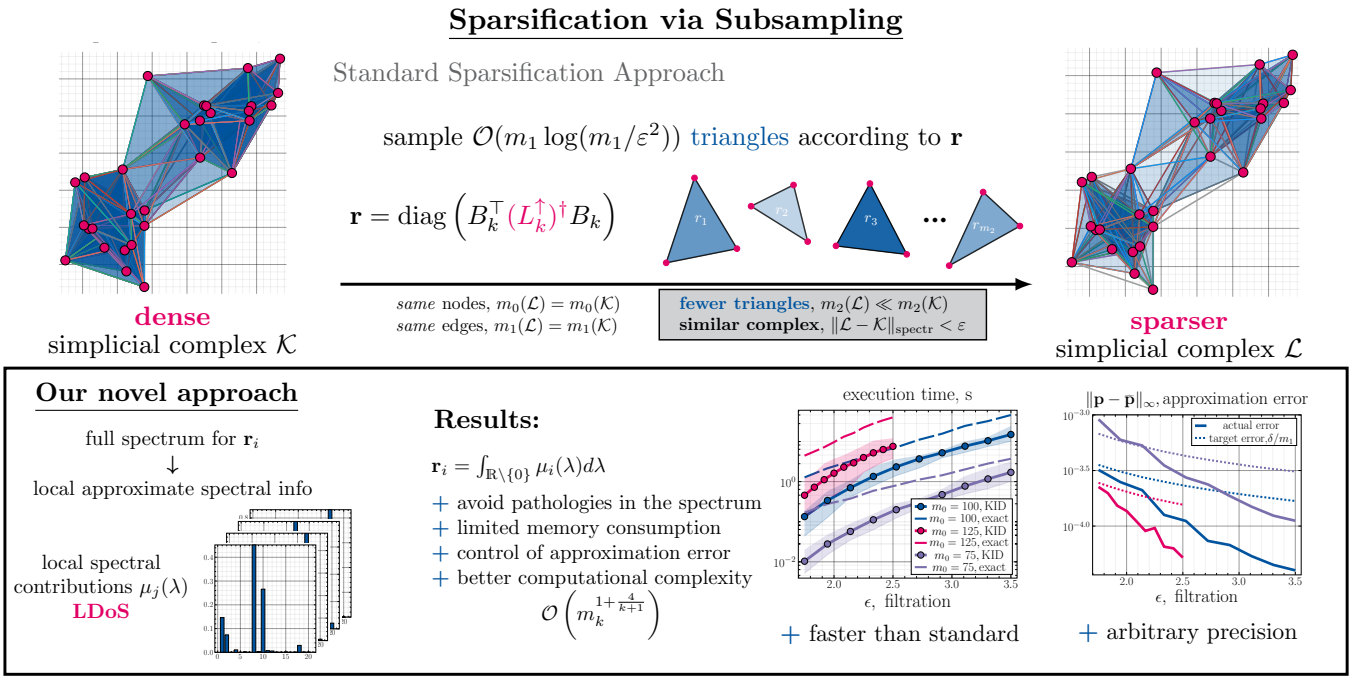

Figure 1: Schematic example of the efficient sparsification of a simplicial complex $\mathcal{K}$ on the level of triangles.

Abstract

Simplicial complexes (SCs), a generalization of graph models for relational data that account for higher-order relations between data items, have become a popular abstraction for analyzing complex data using tools from topological data analysis or topological signal processing. However, the analysis of many real-world datasets leads to dense SCs with a large number of higher-order interactions. Unfortunately, analyzing such large SCs often has a prohibitive cost in terms of computation time and memory consumption. The sparsification of such complexes, i.e., the approximation of an original SC with a sparser simplicial complex with only a log-linear number of high-order simplices while maintaining a spectrum close to the original SC, is of broad interest.

In this work, we develop a novel method for a probabilistic sparsifaction of SCs. At its core lies the efficient computation of sparsifying sampling probability through local densities of states as functional descriptors of the spectral information. To avoid pathological structures in the spectrum of the corresponding Hodge Laplacian operators, we suggest a "kernel-ignoring" decomposition for approximating the sampling probability; additionally, we exploit error estimates to show asymptotically prevailing algorithmic complexity of the developed method. The performance of the framework is demonstrated on the family of Vietoris–Rips filtered simplicial complexes.

Keywords

density of states, simplicial complexes, sparsification, generalized effective resistance

1 Introduction

Graph models for relational data are ubiquitous, enabling the integration of structural information into various tasks such as link prediction, node importance ranking, label propagation, and machine learning algorithms, with numerous applications across disciplines. At the same time, such models are inherently restricted to pairwise interactions between agents, [3, 5]. Many natural systems include polyadic interactions, such as chemical reactions, co-authorship networks, and social connections. To address this limitation, higher-order models of relational data have been introduced, including hypergraphs, motifs, cell and simplicial complexes. Such higher-order models have been shown to influence a variety of processes, such as promoting synchronization [13], impacting label spreading and critical node identification, [26, 36], and supporting higher-order random walks and trajectory classification [31].

However, as the size of a system grows, the number of interactions scales accordingly, rendering most existing network-based algorithms computationally prohibitive. The tractability of the system and the memory consumption are inevitably affected in the same way; the transition to higher-order models only accentuates the problem. As a result, it is natural to posit the question of sparsification: for a given high-order model $\mathcal{K}$, can one find a model $\mathcal{L}$ with similar key properties (such as similar topology or comparable rates of information propagation) but with asymptotically significantly fewer order-k interactions?

This idea of the sparsification is closely related to the Lottery Ticket theorem for the graph neural network (or simplicial complex neural networks in particular, [11, 38]): the transition to a sparser, but similar simplicial complex inside the convolutional layer implies pre-training pruning of connections of each message passing filter.

In the current work, we consider the task of efficient sparsification of simplicial complexes at the level of simplices of order k. For this, the spectral information of the induced higher-order Laplacian operators L_k is critical, as it describes the topology of the complex, governs its stability, may be used to define clusters, and affects the performance of many numerical methods for the analysis of simplicial complexes; consequently, we are computing sparsifiers that maintain the spectrum or specific spectral properties of the original operator L_k (note that the spectral properties of all lower order Laplacians are unaffected by the sparsification). To be more specific, each L_k operator is composed of down- and up-Laplacian terms, $L_k = L_k^{\downarrow} + L_k^{\uparrow}$; since $L_k^{\uparrow}$ desribes up-adjacency between simplices of orders k and $k+1$, one uses its spectrum to control the similarity of the sparsifier. The result by Spielman and Srivastava, [32, 33] and its more recent generalization to simplicial complexes [25] state that for any simplicial complex $\mathcal{K}$ one can find a spectrally close sparser complex $\mathcal{L}$ with $m_{k+1}(\mathcal{L}) = O(m_k(\mathcal{K}) \log m_k(\mathcal{K}))$ where $m_k(\mathcal{K})$ denotes the number of simplices of order k in the simplicial complex $\mathcal{K}$. The sparsifier $\mathcal{L}$ can be sampled from the original complex $\mathcal{K}$ according to a sampling probability proportional to the so-called generalized effective resistance (GER) vector $\mathbf{r}$.

Although a variety of methods for the efficient computation of the GER vector have been proposed for the classical graph case $k = 0$, [9, 19, 34], this remains a major computational bottleneck for higher-order cases. In this work, we show that GER vectors can be directly computed using functional descriptors of the spectral information known as the network's local densities of states (LDoS), introduced in [10]. Due to the Hodge decomposition, [22, 31], high-order up-Laplacian operators $L_k^{\uparrow}$ have high-dimensional kernels which are detrimental for existing quasi-polynomial methods to approximate LDoS. To resolve this issue, we suggest a novel method based on a kernel-ignoring decomposition. Additionally, we provide error estimates, which allows us to guarantee the method's advantageous computational complexity $O\left(\delta^{-3} m_{k+1}^4 m_k^{-3}\right)$ where δ controls the approximation error. The performance of the developed method is illustrated on the family of Vietoris–Rips simplicial complexes, [17], for various density levels and orders of simplices.

Main Contributions . Our main contributions are as follows: **(i)** We show that the sparsification measure used for simplicial complexes is directly related to the local densities of states of a higher-order down-Laplacian $L_{k+1}^{\downarrow}$(whose spectrum is inherited from $L_k^{\uparrow}$, [16]). This measure has previously been defined in terms of the full spectrum of the up-Laplacian; the transition to LDoS enables an efficient approximation. **(ii)** We develop a novel kernel-ignoring decomposition (KID) for the efficient approximation of LDoS; the suggested method avoids pathological spectral structures of the up-Laplacian, which prevented a successful application of the preexisting methods. **(iii)** Finally, we prove that our method outperforms previous approaches in terms of algorithmic complexity, bounding the number of parameters required to obtain a desired approximation error.

Related work

Simplicial complexes. Simplicial complexes generalize models for relational data to higher-order interactions, aligning with the intrinsic topology of the data [24, 27]. The corresponding higher-order Hodge Laplacian operators L_k define homology groups which, for example, describe k-dimensional holes in the complex [22]. In the thermodynamic limit, key topological features are preserved, as L_k operators on simplicial complex induced by point clouds on manifolds converge to their continuous counterparts [6, 7]. The spectral information of Hodge Laplacians has diverse applications: it can be used to determine topological stability of the simplicial complex [16], be exploited for spectral clustering, [12, 15], and trajectory classification through simplicial random walks, [14, 31]. Moreover, L_k can act as structural shift operators for signal processing, [2], and, thus, injected into graph neural networks to model higher-order interactions, [11, 28, 38].

Sparsification. Numerous graph sparsification algorithms have been proposed over the years, each aiming to preserve specific properties of the original system, such as cut costs, [1, 4], clustering, [29], or classification scores, [21]. Due to the prevalence of the graph spectral methods, [32, 33] introduced the concept of spectral sparsification, which preserves the spectrum of the corresponding graph Laplacian operator by sampling edges based on their generalized effective resistance, [35]. This idea was later extended to simplicial complexes in [25]. Several methods have been proposed in the classical graph case to efficiently compute (or approximate) the effective resistance, which is inherently computationally prohibitive as it requires access to the entries of the pseudo-inverse of the Laplacian [9, 34]. This computational challenge directly links

the graph and simplicial sparsification problems to efficient solvers for linear systems involving L_k operators. Unlike the graph case, available approaches for higher-order Laplacians[8, 20, 30] are either uniquely applicable for sparse simplicial complexes or their performance suffers in the densest cases.

Network density of states. The notion of spectral densities as functional descriptors of the spectral information for various operators is well-established across various areas of physics, [37]. It was first introduced in the context of network analysis in [10] where it was observed the presence of pathological structures in the spectrum of the graph Laplacian associated with over-represented motifs in the graph. These structures create potential complications that hinder the application of the kernel polynomial method [37], the algorithm of choice when aiming to approximate the spectral density.

Outline. The rest of the paper is organized as follows: Section 2 provides a brief introduction to simplicial complexes. In Section 3, we present our main sparsification result and its connection to the local spectral density of states, used as efficiently computable proxies for spectral information. Section 4 outlines the proposed novel approach for efficiently computing the sparsifying probability measure. Finally, numerical experiments are presented in Section 5, followed by concluding remarks in Section 6.

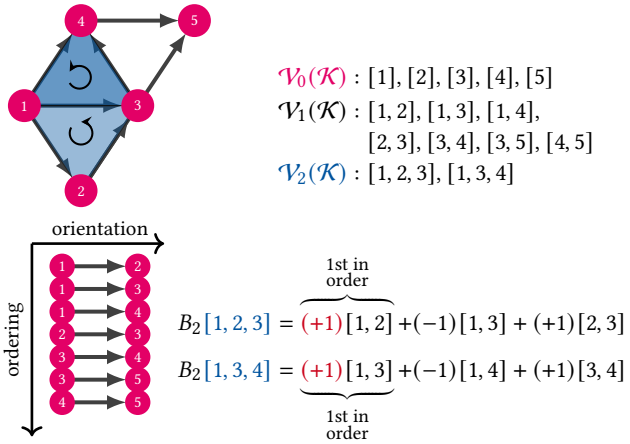

Figure 2: Example of a simplicial complex with ordering and orientation: nodes from $\mathcal{V}_0(\mathcal{K})$ in magenta, triangles from $\mathcal{V}_2(\mathcal{K})$ in blue. Orientation of edges and triangles is shown by arrows; the action of the B_2 operator is given for both triangles. Adapted from [30].

2 Preliminaries

2.1 Notation

We use $\sigma(A)$ to denote the spectrum of a symmetric operator A. $\sigma_+(A)$ denotes the strictly positive part of the spectrum. We say that two symmetric operators A and B are semi-positive ordered $A \succeq B$ iff $A - B \succeq 0$, meaning $A - B$ is a symmetric semi-positive definite operator. Two operators A and B are spectrally ε-close, $A \underset{\varepsilon}{\approx} B$ iff

$$(1-\varepsilon)B \preceq A \preceq (1+\varepsilon)B.$$

We use $\odot$ to denote element-wise matrix multiplication. Finally, for a finite set $\mathcal{S}$, $|\mathcal{S}|$ corresponds to its cardinality.

2.2 Simplicial complexes

Classical pairwise graph models consist of two sets, $G = \{\mathcal{V}_0, \mathcal{V}_1\}$, where $\mathcal{V}_0$ is a set of nodes and $\mathcal{V}_1$ is a set of edges between the nodes. Instead, one may consider the structured generalization, which includes higher-order interactions between the nodes, known as *simplicial complexes*. Let us assume that $\mathcal{V}_0 = \{v_1, v_2, \dots v_{m_0}\}$; then each subset σ of $\mathcal{V}_0$, $\sigma = [v_{i_0}, v_{i_1}, \dots v_{i_k}]$, is called a *simplex* of order k (k-simplex) with its maximal proper subsets of order $(k-1)$ known as *faces* of σ.

Definition 2.1 (Simplicial complex, [22]). A collection of simplices $\mathcal{K}$ on the node set $\mathcal{V}_0$ is a *simplicial complex*, if each simplex σ enters $\mathcal{K}$ with all its faces. Additionally, we say that $\mathcal{K} = \bigcup_{k=0}^{\dim K} \mathcal{V}_k(\mathcal{K})$ where $\mathcal{V}_k(\mathcal{K})$ is a set of simplices of order k and $\dim K$ is the maximal order of simplices in $\mathcal{K}$. We provide a small example of a simplicial complex in Figure 2.

Let us denote the number of k-simplices in $\mathcal{K}$ by m_k, $m_k = |\mathcal{V}_k(\mathcal{K})|$. The sparsity of $\mathcal{K}$ at the level of k-simplices is defined by the relation between m_k and m_{k+1}; we refer to it as the k-sparsity of the simplicial complex. In particular, the 0-sparsity is given by the relation between the number of nodes m_0 and the number of edges m_1 and is what usually defines the sparsity of the graph. Note that the k-sparsity is not on its own indicative of the $(k+1)$-sparsity as the two depend on the intrinsic topology of the simplicial complex. Consequently, if one aims to find a similar but sparser simplicial complex, it is natural to consider such a problem for a fixed k rather than attempting to define a unified notion of sparsifier across all simplex orders.

We now formalize the concept of closeness between simplicial complexes, aiming to identify a simplicial complex $\mathcal{L}$ that is k-sparser than $\mathcal{K}$ while still preserving target properties and structures of the original $\mathcal{K}$. Following [25, 32], we use the notion of sparsification via the spectral closeness of higher-order Laplacian operators L_k that naturally describe the topology of the underlying simplicial complex. For this purpose, the operators L_k are formally defined below.

2.3 Laplacian Operators and Topology

Since each simplex σ enters $\mathcal{K}$ with all its faces, there exists a map matching it to its boundary formed by its faces.

In order to formally define such a map, assume that simplices in $\mathcal{V}_k(\mathcal{K})$ have a fixed ordering and chosen orientations as a matter of bookkeeping (e.g. lexicographically). Then, let us consider linear spaces C_k of formal sums of simplices in $\mathcal{V}_k(\mathcal{K})$; i.e. C_0 is the space of node states, C_1 the space of edge flows, and so on; note that each such space is isomorphic to $\mathbb{R}^{m_k}$, $C_k \cong \mathbb{R}^{m_k}$, and that it can be viewed as the space of simplicial flows on $\mathcal{V}_k(\mathcal{K})$. Then, the action of the boundary map B_k on $\sigma = [v_{i_0}, v_{i_1}, \dots v_{i_k}] \in \mathcal{V}_k(\mathcal{K})$ is defined as an alternating sum:

$$B_k : C_k \mapsto C_{k-1}, \quad B_k\sigma = \sum_{j=0}^{k} (-1)^j \sigma_{\bar{j}}$$

where $\sigma_{\overline{j}}$ denotes the face of σ that does not include v_{i_j}. Given a fixed order and orientation, simplices in $\mathcal{V}_k(\mathcal{K})$ and $\mathcal{V}_{k-1}(\mathcal{K})$ form canonical bases in C_k and C_{k-1} respectively, thus allowing us to use a matrix representation of the boundary operators. For simplicity, from now on, we will always assume this is the case. With a slight abuse of notation, will use the symbol B_k to denote the matrix representation of the boundary operator. An example for a simplicial complex and the action of the boundary operator is provided in Figure 2.

The alternating sum in the definition of boundary operators B_k upholds the fundamental lemma of homology ("a boundary of a boundary is zero"), $B_k B_{k+1} = 0$, and thus induces the topologically sound decomposition of simplicial flows known as *Hodge decompositon*:

$$\mathbb{R}^{m_k} = \overbrace{\operatorname{im} B_k^\top \oplus \underbrace{\ker\left(B_k^\top B_k + B_{k+1}B_{k+1}^\top\right)}_{\ker B_k}}^{\ker B_{k+1}^\top} \oplus \operatorname{im} B_{k+1}. \qquad (1)$$

Definition 2.2 (Hodge Laplacian operator). The operator $L_k = B_k^\top B_k + B_{k+1}B_{k+1}^\top$ is known as the *Hodge* or *higher-order* Laplacian operator and has the following properties:

(1) the elements of $\ker L_k$ in the Hodge decomposition correspond to the k-dimensional holes in $\mathcal{K}$ (connected components for $k = 0$, 1-dimensional holes for $k = 1$, and so on);
(2) the first term $L_k^\downarrow = B_k^\top B_k$ is known as the *down-Laplacian* and describes the relation between k- and $(k-1)$-simplices. For $k = 1$, $\operatorname{im} L_k^\downarrow = \operatorname{im} B_k^\top$ contains so-called gradient flows on the edges;
(3) the second term $L_k^\uparrow = B_{k+1}B_{k+1}^\top$ is known as the *up-Laplacian* and describes the relation between k- and $(k+1)$-simplices. For $k = 1$, $\operatorname{im} L_k^\uparrow = \operatorname{im} B_{k+1}$ contains so-called curl flows.

Higher-order Laplacian operators L_k, as well as their down- and up-terms, are frequently used to describe dynamical processes and random walks on simplicial complexes, [31], to inject simplicial structure into graph neural networks, [11]. Moreover, the eigenvectors of Hodge Laplacians can be used for spectral clustering, [12, 15], or to characterize the stability of topological features, [16].

Note that one can generalize boundary and Laplacian operators to the weighted case without significant difficulty. Let W_k be a diagonal matrix such that its i-th diagonal entry contains the weight of the i-th simplex in $\mathcal{V}_k(\mathcal{K})$, $[W_k]_{ii} = \sqrt{w_k(\sigma_i)}$. Then, $W_{k-1}^{-1} B_k W_k$ provides a weighted version of B_k that preserves all the fundamental topological features of B_k, and the same is true for the corresponding weighted higher-order Laplacian [16].

By its definition, the up-Laplacian $L_k^\uparrow$ describes the relationship between simplicies in $\mathcal{V}_k(\mathcal{K})$ and $\mathcal{V}_{k+1}(\mathcal{K})$. At the same time, its spectrum encodes information about the overall topology of the simplicial complex. As a result, one may use the operator $L_k^\uparrow$ to measure the closeness between simplicial complexes for the task of k-sparsification, as we describe below.

3 Sparsification of Simplicial Complexes

We assume from now on that we consider the k-spasification of the weighted simplicial complex $\mathcal{K}$ and will sometimes omit the index k; to simplify the notations, we assume that $W_{k-1} = I$, though every statement below holds in the general case.

The task of spectral sparsification we are concerned with thus can be generally described as follows:

Problem 1. *For a given weighted simplicial complex $\mathcal{K}$ and the sensitivity level $\varepsilon > 0$, find a simplicial complex $\mathcal{L}$ such that $\mathcal{V}_i(\mathcal{K}) = \mathcal{V}_i(\mathcal{L})$ for all $i = 0, \dots, k$, $\mathcal{V}_{k+1}(\mathcal{L}) \subset \mathcal{V}_{k+1}(\mathcal{K})$ with $m_{k+1}(\mathcal{L}) \ll m_{k+1}(\mathcal{K})$, and the corresponding up-Laplacians are spectrally ε-close, $L_k^\uparrow(\mathcal{L}) \underset{\varepsilon}{\approx} L_k^\uparrow(\mathcal{K})$.*

The seminal result by Spielman and Srivastava, [32], (with generalization into the simplicial case in [25]) suggests that the sparsifier $\mathcal{L}$ can be randomly sampled from $\mathcal{V}_{k+1}(\mathcal{K})$ with a probability measure defined by the generalized effective resistance $\mathbf{r}$ of simplices in $\mathcal{V}_{k+1}(\mathcal{K})$. Specifically,

Theorem 3.1 (Simplicial Sparsification, [25, 32]). *For any $\varepsilon > \frac{1}{\sqrt{m_k}} > 0$, a k-sparse complex $\mathcal{L}$ can be sampled as follows:*

(1) *compute the probability measure $\mathbf{p}$ on $\mathcal{V}_{k+1}(\mathcal{K})$ proportional to the generalized effective resistance (GER) vector $\mathbf{r}$, $\mathbf{p} \sim W_{k+1}^2 \mathbf{r}$, where $\mathbf{r} = \operatorname{diag}\left(B_{k+1}^\top (L_k^\uparrow)^\dagger B_{k+1}\right)$;*
(2) *sample $q(m_k)$ simplices σ_i from $\mathcal{V}_{k+1}(\mathcal{K})$ according to the probability measure $\mathbf{p}$ with $q(m_k) \geq 9C^2 m_k \log(m_k/\varepsilon^2)$, for some absolute constant $C > 0$;*
(3) *form a sparse simplicial complex $\mathcal{L}$ with all the sampled $(k+1)$-simplexes (and all their faces) with the weight $\frac{w_{k+1}(\sigma_i)}{q(m_k)\mathbf{p}(\sigma_i)}$; weights of simplices, repeated during sampling, are accumulated.*

Then, with probability at least $1/2$, the up-Laplacian of the sparsifier $\mathcal{L}$ is spectrally ε-close to the original one, $L_k^\uparrow(\mathcal{L}) \underset{\varepsilon}{\approx} L_k^\uparrow(\mathcal{K})$.

The bottleneck in the sparsification by sampling described in Theorem 3.1 is the computation of the GER vector $\mathbf{r} = \operatorname{diag}\left(B_{k+1}^\top (L_k^\uparrow)^\dagger B_{k+1}\right)$, which requires the pseudo-inverse $(L_k^\uparrow)^\dagger$, a highly computationally demanding operation with complexity $O(m_k^3 + 2k m_k m_{k+1})$. This poses the central problem of the current work:

Problem 2. *Find a computationally efficient and arbitrarily close approximation of the generalized effective resistance vector $\mathbf{r}$ for a given weighted simplicial complex $\mathcal{K}$ and simplices of fixed k-th order.*

In the rest of the paper, we formulate and discuss a novel method for approximating the GER vector using efficiently computable spectral densities of the up-Laplacian operators $L_k^\uparrow$. Note that the sparsification process described in Theorem 3.1 allows for approximate sparsifying measures and exhibits low sensitivity to them, as we discuss below.

Remark 3.2. Since $\mathbf{p}$ is a probability measure over $\mathcal{V}_{k+1}(\mathcal{K})$, it is natural to measure its perturbations in terms of $\frac{1}{m_{k+1}}$, i.e., $\left|\mathbf{p}(\sigma_i) - \mathbf{p}^{(\delta)}(\sigma_i)\right| < \frac{\delta}{m_{k+1}}$ (the size of the perturbation is meaningful only in relation to the actual support of the measure). As shown in Figure 3, random perturbations of $\mathbf{p}$ can, on average, slow down

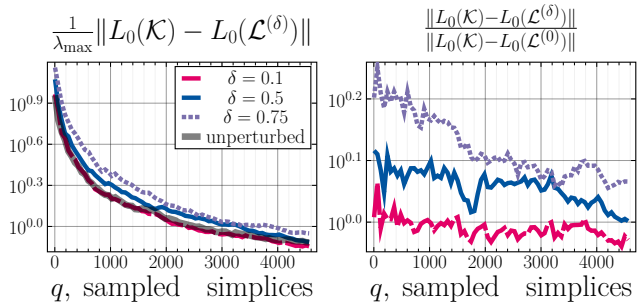

Figure 3: Convergence of the sampled simplicial complex $\mathcal{L}$ to the original complex $\mathcal{K}$ at 0-order in terms of the spectrum of L_0 Laplacian operator. Left pane: convergence rate vs the number of sampled edges q for various perturbed measures $\mathbf{p}^{(\delta)}$. Right pane: convergence rate for chosen values of δ in relation to the unperturbed sparsifier. $m_0 = 100$, $m_1 = \frac{m_0(m_0-1)}{3}$. All curves are averaged over 25 random perturbations for VR-complex (see Section 5).

the convergence of the approximately sampled complex $\mathcal{L}^{(\delta)}$ to the original simplicial complex $\mathcal{K}$ in terms of the number of sampled simplices. However, even for moderately high values of δ, such as $\delta = 0.5$, the convergence rate remains largely unaffected.

4 Sparsification measure via Kernel Ignoring Decomposition of Local Densities of States

In the previous sections, we have outlined the importance of the spectral information of $L_k^{\uparrow}$ and the way it is exploited to obtain the sparsifier $\mathcal{L}$ with the computationally demanding GER vector $\mathbf{r}$. Here we demonstrate that $\mathbf{r}$ can be efficiently approximated via functional descriptors of the up-Laplacian spectrum known as spectral densities or densities of states $\mu_k(\lambda \mid L_k)$, [5]; later we propose a novel method for computing $\mu_k(\lambda \mid L_k)$ that avoids pathological structures in the up- and down-Laplacians' spectra.

4.1 Density of States

Definition 4.1 (Density of States). For a given symmetric matrix $A = Q\Lambda Q^\top$ with $Q^\top Q = I$ and diagonal $\Lambda = \operatorname{diag}(\lambda_1, \dots \lambda_n)$, the *spectral density* or *density of states* (DoS) is defined as

$$\mu(\lambda \mid A) = \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{\delta}(\lambda - \lambda_i)$$

where $\boldsymbol{\delta}(\lambda)$ is a Dirac delta function. Additionally, let $\mathbf{q}_i$ be a corresponding unit eigenvector of A (such that $A\mathbf{q}_i = \lambda_i \mathbf{q}_i$ and $Q = (\mathbf{q}_1 \mid \mathbf{q}_2 \mid \cdots \mid \mathbf{q}_n)$); then one can define a set of local densities of states (*LDoS*):

$$\mu_j(\lambda \mid A) = \sum_{i=1}^{n} \left|\mathbf{e}_j^\top \mathbf{q}_i\right|^2 \boldsymbol{\delta}(\lambda - \lambda_i)$$

with $\mathbf{e}_j$ being the j-th canonical basis vector.

Here, the DoS function $\mu(\lambda \mid A)$ contains the overall spectrum of the operator A while the family of LDoS $\mu_j(\lambda \mid A)$ describes the contribution of the simplex $\sigma_j \in \mathcal{V}_k(\mathcal{K})$ to the spectral information.

Finally, one should note that by definition, DoS and LDoS are generalized functions; hence, the quality of their computation is difficult to assess directly. To this end, one can instead consider their histogram representations:

$$h_i = \int_{x_i}^{x_i+\Delta_x} \mu(\lambda \mid A)d\lambda, \qquad h_i^{(j)} = \int_{x_i}^{x_i+\Delta_x} \mu_j(\lambda \mid A)d\lambda$$

which correspond to the discretized output of the convolution of spectral densities with a smooth approximation of the identity function $K_{\Delta h}$, $[\mu(\lambda \mid A) * K_{\Delta h}]$.

With the notion of spectral density, we now obtain the following reformulation of the GER vector $\mathbf{r}$ and its computation:

Theorem 4.2 (GER through LDoS). *For a given simplicial complex $\mathcal{K}$, with k-th order up-Laplacian $L_k^{\uparrow} = B_{k+1}W_{k+1}^2 B_{k+1}^\top$, the generalized effective resistance $\mathbf{r}$ can be computed through a family of local densities of states $\{\mu_i(\lambda \mid L_{k+1}^{\downarrow})\}$ as follows:*

$$\mathbf{r}_i = \int_{\mathbb{R}} (1 - \mathbb{1}_0(\lambda))\mu_i(\lambda \mid L_{k+1}^{\downarrow})d\lambda$$

Theorem 4.2 implies that it is sufficient to obtain the LDoS family $\{\mu_i(\lambda \mid L_{k+1}^{\downarrow})\}$ for the next down-Laplacian $L_{k+1}^{\downarrow}$ *efficiently* in order to compute the sparsifying probability measure at the level of k-simplices. At the same time, by its definition, the spectral density $\{\mu_i(\lambda \mid A)\}$ requires the complete spectral information of the original operator A and, hence, is not immediately computationally beneficial. To avoid this demanding computational overhead, we leverage the functional nature of the LDoS and obtain an efficient approximation of $\{\mu_i(\lambda \mid L_{k+1}^{\downarrow})\}$ via truncated polynomial expansion.

4.2 Efficient approximation of LDoS

Fast approximations of spectral densities are typically based on Kernel Polynomial Methods (KPM) that schematically operate as follows:

(1) shift the operator $A \mapsto H$ so that $\sigma(H) \subseteq [a; b]$;
(2) consider a polynomial basis $T_m(x)$ on $[a; b]$, orthogonal with respect to the weight function $w(x)$; then it holds

$$\mu_j(\lambda \mid H) = \sum_{m=0}^{\infty} d_{mj} w(\lambda) T_m(\lambda)$$

where the coefficients d_{mj} are known as *moments*. In practice, one is interested in a truncated decomposition of the form $\widehat{\mu}_j(\lambda \mid H) = \sum_{m=0}^{M} d_{mj} w(x) T_m(x)$ for some M;
(3) the values of d_{mj} are functions of $T_m(H)$ and can be efficiently sampled via Monte-Carlo methods, exploiting the three-term recurrence of orthogonal polynomial bases, [5]; we review this step in more detail below.

We point out that typical choices for the shift interval and the polynomial basis are $[-1, 1]$ and the Chebyshev polynomials of the first kind, respectively.

The fundamental limitation of KPM is the polynomial nature of the decomposition, which may require a large number of moments M for $\widehat{\mu}_j(\lambda \mid H)$, in order to produce an accurate approximation for "pathologic" functions that are far from being polynomials. In the case of the DoS and LDoS, a particularly pathological setting

is associated with eigenvalues of high multiplicity, which result in dominating "spikes" in the histogram representations of the spectral densities. In this case, one would have to approximate an outlier with a polynomial function. In the case of the classical graph Laplacian L_0 and the adjacency matrix, these spikes may be caused by over-represented motifs in the graph [5]. However, in this setting, one knows the closed form of the corresponding eigenspace, and thus, the over-represented eigenvalues can be explicitly filtered out. For the general case of up- and down-Laplacians $L_k^{\uparrow}$ and $L_k^{\downarrow}$ of order $k > 0$, the eigenspaces with high multiplicity are unavoidable due to the Hodge decomposition, (1): indeed, since $\operatorname{im} B_k^\top \subseteq \ker B_{k+1}^\top = \ker L_k^{\uparrow}$ and $\operatorname{im} B_{k+1} \subseteq \ker B_k = \ker L_k^{\downarrow}$, both operators have large kernels which are detrimental to KPM approximation; moreover, kernel's bases depend on the topology of the simplicial complex and cannot be easily estimated.

Below, we propose a novel modified method for approximating LDoS that intentionally avoids the spike in the kernel of $L_{k+1}^{\downarrow}$, with all the necessary definitions.

4.3 Kernel-ignoring Decomposition

As discussed above, the quality of the KPM approximation of LDoS $\{\mu_j(\lambda \mid L_k^{\downarrow})\}$ suffers from the large null space of the operator; at the same time, Theorem 4.2 suggests that the target GER vector ignores the values of $\mu_j(\lambda \mid L_{k+1}^{\downarrow})$ associated with the kernel due to the $(1 - \mathbb{1}_0(\lambda))$ term. Note that, since $\mu_j(\lambda \mid L_{k+1}^{\downarrow}) \approx \sum_{m=0}^{M} d_{mj} w(\lambda) T_m(\lambda)$ is a functional decomposition, it does not allow local errors; instead, the approximation error associated with the spike at $\lambda = 0$ spreads across the whole domain. For that reason, we suggest a modified shifting technique that leads to a decomposition that ignores the singular value of λ associated with the operator's null space.

In pre-existing methods, one sets $H = \frac{2}{\lambda_{\max}} L_{k+1}^{\downarrow} - I$ where $\lambda_{\max}$ denotes the largest eigenvalue of $L_{k+1}^{\downarrow}$, so that $\sigma(H) \subseteq [-1; 1]$ and the null eigenvectors of $L_{k+1}^{\downarrow}$ are shifted to $-1 \in \sigma(H)$. Instead, we set $H = \frac{1}{\lambda_{\max}} L_{k+1}^{\downarrow}$ and define a symmetrized version of the LDoS as:

$$\tilde{\mu}_j(\lambda \mid H) = \begin{cases} \mu_j(\lambda \mid H), & \text{if } \lambda \in (0, 1] \\ 0, & \text{if } \lambda = 0 \\ -\mu_j(-\lambda \mid H), & \text{if } \lambda \in [-1, 0) \end{cases} \tag{2}$$

Note that the support of the symmetrized $\tilde{\mu}_j(\lambda \mid H)$ still falls within $[-1; 1]$, and the spike associated with the value $\lambda = 0$ is by design tied to 0.

The remainder of the approximation approach is essentially inherited from the KPM method: let us assume that $T_m(x)$ are Chebyhev polynomials of the first kind. Specifically,

$$T_0(x) = 1,\ T_1(x) = x, \qquad T_{m+1}(x) = 2xT_m(x) - T_{m-1}(x) \tag{3}$$

which as an orthonormal basis on $[-1, 1]$ (with respect to the weight function $w(x) = \frac{2}{(1+\delta_{0m})\pi\sqrt{1-x^2}}$).

Since $\tilde{\mu}_j(\lambda \mid H)$ is an odd function by design, we decompose

$$\tilde{\mu}_j(\lambda \mid H) = \sum_{m=0}^{\infty} d_{mj} w(\lambda) T_m(\lambda) \tag{4}$$

where $d_{mj} = 0$ for even m, thus requiring the decomposition by odd Chebyshev polynomials.

Due to the orthonormality of $T_m(x)$ (see the derivation in Appendix B), we have

$$d_{mj} = \left\langle \tilde{\mu}_j(\lambda \mid H), T_m(x) \right\rangle = 2 \sum_{\lambda_i \in \sigma(H) \setminus \{0\}} \left| \mathbf{e}_j^\top \mathbf{q}_i \right|^2 T_m(\lambda_i) = 2 \left[T_m(H) \right]_{jj} \tag{5}$$

thus, $d_{m\bullet} = 2 \operatorname{diag}(T_m(H))$. Note that instead of computing $T_m(H)$ one can use Monte-Carlo estimations for the trace and diagonal, specifically:

$$\operatorname{diag} X = \mathbb{E}\left[\mathbf{z} \odot X\mathbf{z} \right], \qquad \operatorname{diag} X \approx \frac{1}{N_z} (Z \odot XZ)\mathbf{1} \tag{6}$$

where $\mathbf{z}$ is a vector of i.i.d. random variables with zero mean and unit variance, $\mathbf{1}$ is a vector of ones, and Z is a matrix collecting N_z copies of $\mathbf{z}$ column-wise [18, 23].

The MC-sampling reduces $\operatorname{diag}(T_m(H))$ calculations to simple `matvec` operations, which can be efficiently updated for the next order of moments $d_{m+1,\bullet}$ due to the recurrent definition of $T_m(x)$: indeed, assume we store the values of $T_i(H)Z$ for $i = 0 \dots m$; then in order to get $T_{m+1}(H)Z$ one needs to compute $T_{m+1}(H)Z = 2H \cdot (T_m(H)Z) - (T_{m-1}(H)Z)$ requiring only one additional `matvec` operation. As a result, the computational cost of the approximation is fixed to $O(N_z M \operatorname{nnz}(H))$, where $O(\operatorname{nnz}(H))$ is the cost of one `matvec` operation for the operator H.

4.4 Error Propagation and the Choice of Constants

The computational complexity of the KID method described above does not allow for a straightforward comparison with the direct GER vector computation via the pseudo-inverse of $L_k^{\uparrow}$ because it is formulated in terms of the approximation parameters N_z and M rather than the number of simplices m_k. Nonetheless, this can be addressed through error estimates for the LDoS approximation, as we demonstrate below.

Let us consider the histogram representations of the exact symmetrized LDoS $h^{(j)} = \left[\tilde{\mu}_j(\lambda \mid A) * K_{\Delta h} \right]$ and KPM approximation $\widehat{h}_M^{(j)} = \left[\left(\sum_{m=0}^{M} d_{mj} w(\lambda) T_m(\lambda) \right) * K_{\Delta h} \right]$ where the moments d_{mj} are MC-sampled. Using the estimation bound from [5, Thm 4.2], we obtain:

$$\mathbb{E} \left\| h^{(j)} - \widehat{h}_M^{(j)} \right\|_\infty \le \frac{1}{\Delta h} \left(\frac{6L}{M} + \frac{2\|K_{\Delta h}\|_\infty}{\pi\sqrt{N_z}} \right) \tag{7}$$

where L is the Lipschitz constant of $h^{(j)}$. This estimate is the key to obtaining the following (proof moved to the appendix):

Theorem 4.3. *For any fixed $\delta > 0$, let $\mathbf{p}$ and $\widehat{\mathbf{p}}$ be the exact and the KID-approximate sparsifying probability measures for a given sufficiently dense simplicial complex $\mathcal{K}$. If the approximation $\widehat{\mathbf{p}}$ is obtained with $M \ge 24L \frac{m_{k+1}}{\delta m_k}$ and $N_z \ge \frac{8\|K_{\Delta h}\|^2}{\pi^2} \frac{m_{k+1}^2}{\delta^2 m_k^2}$, then*

$$\|\mathbf{p} - \widehat{\mathbf{p}}\|_\infty \le \delta \tag{8}$$

and the computational complexity of the KID approximation is

$$O\left(\delta^{-3} m_{k+1}^4 m_k^{-3} \right).$$

To get a more stream-lined understanding of the overall complexity, we recall that $m_{k+1} = O\left(m_k^{1+\frac{1}{k+1}}\right)$ for a simplicial complex lifted from a completely connected graph, resulting into the worst-case complexity of $O\left(\delta^{-3} m_k^{1+\frac{4}{k+1}}\right)$. In other words, our approximation is not worse than direct computation for $k = 1$ (even in the densest case) and is asymptotically linear in k.

5 Benchmarking

In this section, we show the performance of the KID approximation for the GER vector $\mathbf{r}$ and its computational complexity. In particular, through our experimental evaluation, we aim to do the following:

- **(i)** support the asymptotic estimate for the approximation error (7) in terms of the number of moments M and the number of MC-vectors N_z;
- **(ii)** support the computational complexity of the KID approximation using the (scaled) oracle choice for the parameters, Theorem 4.3;
- **(iii)** compare the actual execution time of the approximation to the direct computation for complexes of different sizes and densities.

Vietoris–Rips filtration. Theorem 4.3 and Equation (7) describe the performance of the developed method in terms of the number of simplices m_k. In order to appropriately numerically illustrate these behaviours, one should consider a family of arbitrarily large and dense simplicial complexes. For this reason, we opt to use simplicial complexes induced by the filtration procedure on point clouds. Formally, we proceed as follows:

(1) we consider m_0 points embedded in $\mathbb{R}^2$, sampled randomly in two clusters, i.e., $\frac{m_0}{2}$ points are sampled from $\mathcal{N}(\mathbf{0}, I)$ and $\frac{m_0}{2}$ points are sampled from $\mathcal{N}(c\mathbf{1}, I)$, for some $c > 0$;
(2) then, for a fixed filtration threshold $\epsilon > 0$, a simplex $\sigma = [v_{i_1}, \ldots v_{i_p}]$ on these nodes enters the generated complex $\mathcal{K}$ if and only if $d_{\mathcal{M}}(v_{i_j}, v_{i_k}) \le \epsilon$ for all paris j and k.

This straightforward filtration is known as Vietoris–Rips filtration, and the corresponding complex $\mathcal{K}$ as a VR-complex. An illustrative example is provided in Figure 4. In the chosen setup, the value of the filtration parameter ϵ naturally governs the density of the generated simplicial complexes of every order, as shown by the right panel in Figure 4: larger values of ϵ define complexes with a higher number of edges, triangles, tetrahedrons, etc., until every possible simplex is included in $\mathcal{K}$.

Parameter choice and computational complexity. The error estimate from Equation 7 suggests that the approximation error for the sparsifying norm $\mathbf{p}$ scales as M^{-1} in terms of the number of moments and as $N_z^{-1/2}$ in terms of the number of Monte-Carlo vectors (MC-vectors). To illustrate this behaviour, we fix one of the parameters (M or N_z) to their theoretical estimates provided by Theorem 4.3 and demonstrate the dynamic of the error $\|\mathbf{p} - \widehat{\mathbf{p}}\|_\infty$ as the function of the other parameter. As shown by Figure 5, the overall scaling law coincides with the estimates of Equation 7 in the case of 1-sparsification for $L_1^{\uparrow}$ operator. Note that all experiments are conducted in the at least minimally-dense setting, i.e. $m_2 \ge m_1 \ln m_1$.

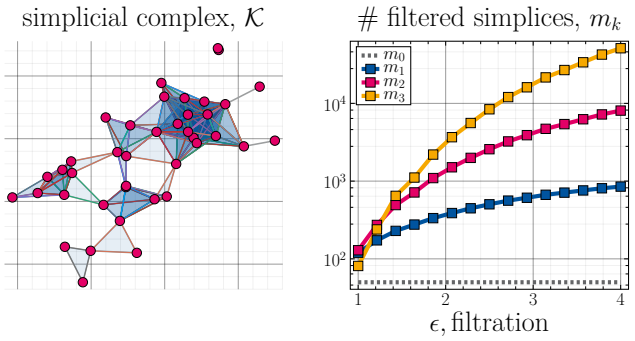

Figure 4: Example of VR-filtration. Left pane: point cloud with $m_0 = 40$ and filtration $\epsilon = 1.5$, inter-cluster distance $c = 3$. Right pane: dynamics of the number of simplices of different orders for varying filtration parameter ϵ.

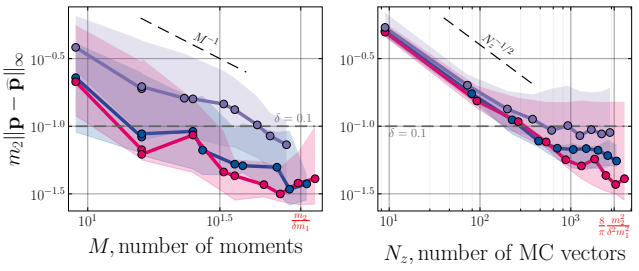

Figure 5: Dependence of the approximation error $\|\mathbf{p} - \widehat{\mathbf{p}}\|_\infty$ on the number of moments M and number of MC vectors N_z. Values are tested up to (scaled) theoretical bounds from Thm 4.3 (in red); line colors correspond to varying m_0 in the point cloud. Left pane: errors vs the number of moments M with fixed theoretical N_z; right pane: errors vs the number of MC vectors N_z with fixed theoretical M. Errors are averaged over several generated VR-complexes; colored areas correspond to the spread of values.

Here, we explicitly highlight two observations from Figure 5: (1) larger and denser simplicial complexes tend to exhibit faster convergence in both parameters (especially in the number of moments M), and (2) Theorem 4.3 provides theoretical (greedy) estimates for M and N_z that are sufficient for achieving the target approximation quality δ and can be interpreted asymptotically. Consequently, one may choose scaled (and empirically sufficient) values for these parameters:

$$M = \left\lceil \frac{m_{k+1}}{\delta\, m_k} \right\rceil \quad \text{and} \quad N_z = \left\lceil \frac{1}{10}\, \frac{8}{\pi^2}\, \frac{m_{k+1}^2}{\delta^2\, m_k^2} \right\rceil.$$

Given this choice of parameters, in Figure 6 we demonstrate that the complexity estimate $O(\frac{m_{k+1}^4}{m_k^3})$ from Theorem 4.3 aligns with the actual execution time of the KID approximation for varying filtration parameters ϵ in the cases of 1- and 2-sparsification of VR-complexes.

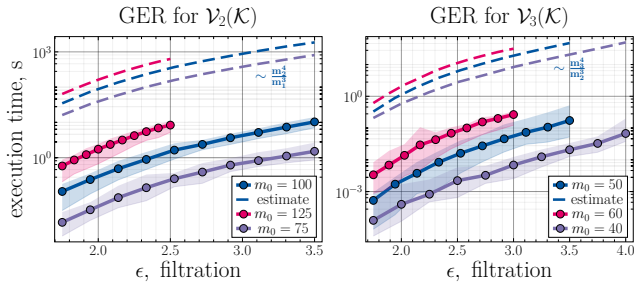

Figure 6: Execution time of KID approximation for effective resistance of triangles, $\mathcal{V}_2(\mathcal{K})$ (left), and tethrahedrons, $\mathcal{V}_3(\mathcal{K})$ (right). Line colors correspond to varying m_0 in the point cloud; theoretical estimation of the computational complexity is given in dash. Execution times are averaged over several generated VR-complexes; colored areas correspond to the spread of values.

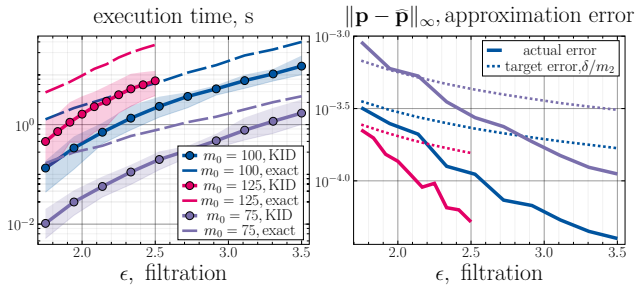

Figure 7: Computation time comparison between KID-approximation $\widehat{\mathbf{p}}$, solid line, and directly computed sparsifying measure $\mathbf{p}$, dashed line (left), and corresponding approximation error $\|\widehat{\mathbf{p}} - \mathbf{p}\|_\infty$ (right). Target approximation error is given in dotted line (right pane); line colors correspond to varying m_0 in the point cloud. Execution times are averaged over several generated VR-complexes.

Comparison with the direct computation. Finally, we compare the execution time of the KID approximation with that of the direct computation of the sparsifying measure $\mathbf{p}$ for 1-sparsification, using the approximation parameters mentioned above (see Figure 7). Note that although the densest case complexity estimate $O(\delta^{-3}\, m_k^{1+\frac{4}{k+1}})$ suggests that the KID method's execution time might be comparable to direct computation, in practice the developed algorithm is significantly faster while still maintaining the target approximation error $\|\mathbf{p} - \widehat{\mathbf{p}}\|_\infty \le \frac{\delta}{m_2}$.

Additionally, we note that the performance of the direct computation of $\mathbf{p}$ for the largest considered point cloud with $m_0 = 125$ highlights another important advantage of the KID approximation: reduced memory consumption. Indeed, whether one uses the definition of the GER vector

$$\mathbf{r} = \operatorname{diag}\left(B_{k+1}^\top (L_k^\uparrow)^\dagger B_{k+1}\right)$$

or the reformulation in terms of the right singular vector from Theorem 4.2, a full SVD of $L_k^\uparrow$ is required. In the case of point clouds with $m_0 = 125$, denser VR-complexes lead to real-valued matrices of size $10^4 \times 10^4$, resulting in substantial memory demands for the SVD. By contrast, the KID approximation avoids this decomposition and restricts the additional memory usage to storing Monte-Carlo matrices Z and their `matvecs` of dimension $m_{k+1} \times N_z$, which is comparatively smaller.

6 Discussion

In this work we have proposed a fast method of approximating generalized effective resistance vector for simplices of an arbitrary order k inside simplicial complex $\mathcal{K}$ with the algorithmic complexity $O\left(\delta^{-3} m_k^{1+\frac{4}{k+1}}\right)$, Theorem 4.3, allowing for efficient k-sparsification of $\mathcal{K}$ through subsampling, Theorem 3.1. The novel approach is based on the connection between GER vector $\mathbf{r}$ and the family of local density of states of the corresponding higher-order down-Laplacian operator $L_{k+1}^\downarrow$, Theorem 4.2. We avoid the pathologic behaviour of the pre-existing KPM approximation methods for LDoS by suggesting a kernel-ignoring decomposition de facto decomposing symmetrized spectral densities via odd Chebyshev polynomials. Developed approach follows the theoretical estimates, Theorem 4.3 and Equation (7), for the approximation error which allows us to choose the number of moments M and number of Monte Carlo vectors N_z controlling the final approximation error which is shown to be sufficient at a moderate level for a close-to-efficient sparsification, Figure 3. Given the fact that the developed method is only dependent on $L_k^\uparrow$ being a Hodge Laplacian, it can also be directly applied to cell complexes.

Given the sufficient improvement of the asymptotic of the computational complexity of the method for $k \ge 1$, the introduction of the sparsified complex in label spreading, spectral clustering or generic simplicial complex GNN tasks may sufficiently decrease priorly prohibitive computational costs, [12, 38]; additionally, the cost of trajectory classification as well as landmark detection algorithms can be directly scaled down by transitioning to a sparser, but spectrally similar model, [14]. Separately, one may notice that the existence of the efficient sparsifier $\mathcal{L}$ effectively breaches the gap between dense complexes and ones for which one can obtain a preconditioner through the collapsible subcomplex, [30], resulting into an efficient solver for linear systems $L_k^\uparrow \mathbf{x} = \mathbf{b}$ for arbitrary simplicial complex. Additionally, the developed method admits graph sparsification as a special case; at the same time, the computational complexity for $k = 0$ would clearly exceed the complexity of preexisting graph sparsification algorthims. Note that we do not suggest the application of KID-approximation for the case of the classical graph, since it is not the focus of the current work.

Finally, one may note, that the approximation quality given in Theorem 4.3 may still be improved in the pathological scenarios such as over-represented motifs in the graph (see [10]) through similar filtration technique. Whilst we avoid explicitly defining a generalized motif for simplices of the arbitrary order and whether their detrimental effect of LDoS is generic or non-generic for simplicial complexes, it is promising future venue of research as well

as other applications of KID-approximated LDoS for simplicial complexes.

Acknowledgments

AS and MTS acknowledge funding by the German Research Council (DFG) within the Collaborative Research Center "Sparsity and Singular Structures" (SfB 1481); MTS acknowledges funding by the European Union (ERC, HIGH-HOPeS, 101039827). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The research of NG was supported by funds from the Italian MUR (Ministero dell'Universitá e della Ricerca) within the PRIN 2022 Project "Advanced numerical methods for time dependent parametric partial differential equations with applications" and the PRO3 joint project entitled "Calcolo scientifico per le scienze naturali, sociali e applicazioni: sviluppo metodologico e tecnologico". FT is partially funded by the PRIN-MUR project MOLE code: 2022ZK5ME7 Principal MUR D.D. financing decree n. 20428 of November 6th, 2024, CUP B53C24006410006. FT and NG were supported from MUR-PRO3 grant STANDS and the PRIN-PNRR project FIN4GEO within the European Union's Next Generation EU framework, Mission 4, Component 2, CUP P2022BNB97. FT and NG are members of INdAM GNCS (Gruppo Nazionale di Calcolo Scientifico).

References

[1] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. 2012. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*. 5–14.

[2] Sergio Barbarossa and Stefania Sardellitti. 2020. Topological signal processing over simplicial complexes. *IEEE Transactions on Signal Processing* 68 (2020), 2992–3007.

[3] Federico Battiston, Giulia Cencetti, Iacopo Iacopini, Vito Latora, Maxime Lucas, Alice Patania, Jean-Gabriel Young, and Giovanni Petri. 2020. Networks beyond pairwise interactions: Structure and dynamics. *Physics reports* 874 (2020), 1–92.

[4] András A Benczúr and David R Karger. 1996. Approximating st minimum cuts in Õ (n 2) time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. 47–55.

[5] Austin R Benson, David F Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.

[6] Yu-Chia Chen and Marina Meila. 2021. The decomposition of the higher-order homology embedding constructed from the k-Laplacian. *Advances in Neural Information Processing Systems* 34 (2021), 15695–15709.

[7] Yu-Chia Chen, Weicheng Wu, Marina Meilă, and Ioannis G Kevrekidis. 2021. Helmholtzian Eigenmap: Topological feature discovery & edge flow learning from point cloud data. *arXiv preprint arXiv:2103.07626* (2021).

[8] Michael B Cohen, Brittany Terese Fasy, Gary L Miller, Amir Nayyeri, Richard Peng, and Noel Walkington. 2014. Solving 1-laplacians in nearly linear time: Collapsing and expanding a topological ball. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 204–216.

[9] Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. 2014. Solving SDD linear systems in nearly m log1/2 n time. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 343–352.

[10] Kun Dong, Austin R Benson, and David Bindel. 2019. Network density of states. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1152–1161.

[11] Stefania Ebli, Michaël Defferrard, and Gard Spreemann. 2020. Simplicial neural networks. *arXiv preprint arXiv:2010.03633* (2020).

[12] Stefania Ebli and Gard Spreemann. 2019. A notion of harmonic clustering in simplicial complexes. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 1083–1090.

[13] Lucia Valentina Gambuzza, Francesca Di Patti, Luca Gallo, Stefano Lepri, Miguel Romance, Regino Criado, Mattia Frasca, Vito Latora, and Stefano Boccaletti. 2021. Stability of synchronization in simplicial complexes. *Nature communications* 12, 1 (2021), 1255.

[14] Vincent P Grande, Josef Hoppe, Florian Frantzen, and Michael T Schaub. 2024. Topological Trajectory Classification and Landmark Inference on Simplicial Complexes. *arXiv preprint arXiv:2412.03145* (2024).

[15] Vincent P Grande and Michael T Schaub. 2024. Node-Level Topological Representation Learning on Point Clouds. *arXiv preprint arXiv:2406.02300* (2024).

[16] Nicola Guglielmi, Anton Savostianov, and Francesco Tudisco. 2023. Quantifying the Structural Stability of Simplicial Homology. *Journal of Scientific Computing* 97, 1 (Aug. 2023), 2.

[17] Jean-Claude Hausmann et al. 1994. *On the Vietoris-Rips complexes and a cohomology theory for metric spaces*. Université de Genève-Section de mathématiques.

[18] Michael F Hutchinson. 1989. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation* 18, 3 (1989), 1059–1076.

[19] Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. 2013. A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 911–920.

[20] Rasmus Kyng and Sushant Sachdeva. 2016. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 573–582.

[21] Jiayu Li, Tianyun Zhang, Hao Tian, Shengmin Jin, Makan Fardad, and Reza Zafarani. 2022. Graph sparsification with graph convolutional networks. *International Journal of Data Science and Analytics* (2022), 1–14.

[22] Lek-Heng Lim. 2020. Hodge Laplacians on graphs. *Siam Review* 62, 3 (2020), 685–715.

[23] Raphael A Meyer, Cameron Musco, Christopher Musco, and David P Woodruff. 2021. Hutch++: Optimal stochastic trace estimation. In *Symposium on Simplicity in Algorithms (SOSA)*. SIAM, 142–155.

[24] James R Munkres. 2018. *Elements of algebraic topology*. CRC press.

[25] Braxton Osting, Sourabh Palande, and Bei Wang. 2022. Spectral sparsification of simplicial complexes for clustering and label propagation. *Journal of computational geometry* 11, 1 (2022).

[26] Konstantin Prokopchik, Austin R Benson, and Francesco Tudisco. 2022. Nonlinear Feature Diffusion on Hypergraphs. In *International Conference on Machine Learning*. PMLR, 17945–17958.

[27] Daniel G Quillen. 2006. *Homotopical algebra*. Vol. 43. Springer.

[28] T Mitchell Roddenberry, Nicholas Glaze, and Santiago Segarra. 2021. Principled simplicial neural networks for trajectory prediction. In *International Conference on Machine Learning*. PMLR, 9020–9029.

[29] Venu Satuluri, Srinivasan Parthasarathy, and Yiye Ruan. 2011. Local graph sparsification for scalable clustering. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. 721–732.

[30] Anton Savostianov, Francesco Tudisco, and Nicola Guglielmi. 2024. Cholesky-like Preconditioner for Hodge Laplacians via Heavy Collapsible Subcomplex. *SIAM J. Matrix Anal. Appl.* 45, 4 (2024), 1827–1849. https://doi.org/10.1137/23M1626396

[31] Michael T Schaub, Austin R Benson, Paul Horn, Gabor Lippner, and Ali Jadbabaie. 2020. Random walks on simplicial complexes and the normalized Hodge 1-Laplacian. *SIAM Rev.* 62, 2 (2020), 353–391.

[32] Daniel A Spielman and Nikhil Srivastava. 2008. Graph sparsification by effective resistances. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. 563–568.

[33] Daniel A Spielman and Shang-Hua Teng. 2011. Spectral sparsification of graphs. *SIAM J. Comput.* 40, 4 (2011), 981–1025.

[34] Daniel A Spielman and Shang-Hua Teng. 2014. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM J. Matrix Anal. Appl.* 35, 3 (2014), 835–885.

[35] Prasad Tetali. 1991. Random walks and the effective resistance of networks. *Journal of Theoretical Probability* 4 (1991), 101–109.

[36] Francesco Tudisco, Austin R Benson, and Konstantin Prokopchik. 2021. Nonlinear higher-order label spreading. In *Proceedings of the Web Conference 2021*. 2402–2413.

[37] Alexander Weiße, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. 2006. The kernel polynomial method. *Reviews of modern physics* 78, 1 (2006), 275–306.

[38] Maosheng Yang, Elvin Isufi, and Geert Leus. 2022. Simplicial convolutional neural networks. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8847–8851.

A Proof of Theorem 4.2

Proof. Let $B_{k+1}W_{k+1} = USV^\top$ where S is diagonal and invertible and both U and V are orthogonal (so it is a truncated SVD

decomposition of $B_{k+1}W_{k+1}$ matrix with eliminated obsolete kernel). Then:

$$(L_k^{\uparrow})^{\dagger} = \left(B_{k+1}W_{k+1}^2B_{k+1}^{\top}\right)^{\dagger} = \left(US^2U^{\top}\right)^{\dagger} = US^{-2}U^{\top} \tag{9}$$

$$\begin{aligned}\mathbf{r} &= \operatorname{diag}\left(W_{k+1}B_{k+1}^{\top}(L_k^{\uparrow})^{\dagger}B_{k+1}W_{k+1}\right) = \\ &= \operatorname{diag}\left(VSU^{\top}US^{-2}U^{\top}USV^{\top}\right) = \operatorname{diag}\left(VV^{\top}\right)\end{aligned} \tag{10}$$

As a result, $\mathbf{r}_i = \|V_{i\cdot}\|^2 = \sum_j |v_{ij}|^2$, so the i-th entry of the resistance is defined by the sum of square of i-th components of eigenvectors $\mathbf{v}_j$ of $L_{k+1}^{\downarrow} = W_{k+1}B_{k+1}^{\top}B_{k+1}W_{k+1}$ operator where $\mathbf{v}_j \perp \ker L_{k+1}^{\downarrow}$.

Note that

$$\mu_i(\lambda \mid L_{k+1}^{\downarrow}) = \sum_{j=1}^{m_{k+1}} \left|\mathbf{e}_i^{\top}\mathbf{q}_j\right|^2 \boldsymbol{\delta}\left(\lambda-\lambda_j\right) = \sum_{j=1}^{m_{k+1}} \left|q_{ij}\right|^2 \boldsymbol{\delta}\left(\lambda-\lambda_j\right) \tag{11}$$

so

$$\begin{aligned}\mathbf{r}_i = \|V_{i\cdot}\|^2 = \sum_j |v_{ij}|^2 = \int_{\mathbb{R}\setminus\{0\}} \sum_{j=1}^{m_{k+1}} \left|q_{ij}\right|^2 \boldsymbol{\delta}\left(\lambda-\lambda_j\right) d\lambda = \\ = \int_{\mathbb{R}\setminus\{0\}} \mu_i(\lambda \mid L_{k+1}^{\downarrow}) d\lambda = \int_{\mathbb{R}} (1-\mathbb{1}_0(\lambda))\mu_i(\lambda \mid L_{k+1}^{\downarrow}) d\lambda\end{aligned} \tag{12}$$

□

B Derivation for Equation 5

Let $\mu_j(x \mid A)$ be a general case for LDoS; then for an arbitrary polynomial function $f(x)$ and $A = Q\Lambda Q^{\top}$,

$$\begin{aligned}\left\langle \mu_j(x \mid A), f(x)\right\rangle &= \int_{-\infty}^{+\infty} \boldsymbol{\delta}\left(x-\lambda_i\right) f(x) dx = \\ &= \sum_{i=1}^{n} \left|\mathbf{e}_j^{\top}\mathbf{q}_i\right|^2 \int_{-\infty}^{+\infty} \boldsymbol{\delta}\left(x-\lambda_i\right) f(x) dx = \\ &= \sum_{i=1}^{n} \left|\mathbf{e}_j^{\top}\mathbf{q}_i\right|^2 f(\lambda_i) = \sum_{i=1}^{n} \left|q_{ji}\right|^2 f(\lambda_i)\end{aligned} \tag{13}$$

Since $f(A) = Qf(\Lambda)Q^{\top}$, we get

$$[f(A)]_{jj} = \mathbf{e}_j^{\top}Qf(\Lambda)Q^{\top}\mathbf{e}_j = \mathbf{q}_j^{\top}f(\Lambda)\mathbf{q}_j = \sum_{i=1}^{n} \left|q_{ji}\right|^2 f(\lambda_i) \tag{14}$$

The case of symmetrized $\tilde{\mu}_j(x \mid H)$ is marginally different; indeed, since $T_m(x)$ is odd and $T_m(0) = 0$:

$$\begin{aligned}\left\langle \tilde{\mu}_j(x \mid H), T_m(x)\right\rangle &= \sum_{\substack{i=1\\ \lambda_i\neq 0}}^{n} \left|q_{ji}\right|^2 \left(T_m(\lambda_i) - T_m(-\lambda_i)\right) = \\ &= 2\sum_{\substack{i=1\\ \lambda_i\neq 0}}^{n} \left|q_{ji}\right|^2 T_m(\lambda_i) = 2\sum_{i=1}^{n} \left|q_{ji}\right|^2 T_m(\lambda_i)\end{aligned} \tag{15}$$

So

$$d_{mj} = 2\sum_{i=1}^{n} \left|q_{ji}\right|^2 T_m(\lambda_i) = 2\left[T_m(H)\right]_{jj} \tag{16}$$

C Proof of Theorem 4.3

Before showing the result directly, we briefly show an auxiliary fact:

Lemma C.1. *For a given simplicial complex $\mathcal{K}$ and GER vector $\mathbf{r}$ it holds that $\|\mathbf{r}\|_1 = m_k - \sum_{i=-1}^{k-1}(-1)^{k-1-i}(m_i - \beta_{i+1})$ where $\beta_k = \dim\ker L_k$ denotes k-th Betti's number and $m_{-1} = 0$.*

Proof. Note that by the proof of Theorem 4.2, $\mathbf{r}_i \geq 0$, and

$$\|\mathbf{r}\|_1 = \sum_i \mathbf{r}_i = \sum_{i,j} |v_{ij}|^2 \tag{17}$$

Since each of the right singular vectors $V_{\cdot j}$ of $B_{k+1}W_{k+1}$ has unit length, then $\sum_{i,j}|v_{ij}|^2 =$ number of columns of $V = m_k - \dim\ker L_k^{\uparrow}$ given the truncated SVD. Given Hodge decomposition (1), $\dim\ker L_k^{\uparrow} = \dim\operatorname{im} B_k^{\top} + \dim\ker L_k = \sum_{i=-1}^{k-1}(-1)^{k-1-i}(m_i - \beta_{i+1})$ due to the spectral inheritance principle for Hodge Laplacians, [16], yielding the result. □

Finally, to show the estimate for the approximation of the sparsifying norm, we consider the estimate

$$\mathbb{E}\left\|h^{(j)} - \widehat{h}_M^{(j)}\right\|_{\infty} \leq \frac{1}{\Delta h}\left(\frac{6L}{M} + \frac{2\|K_{\Delta h}\|_{\infty}}{\pi\sqrt{N_z}}\right) \tag{18}$$

which verbatim translates to the estimate on GER vector $\mathbf{r}$. Given each $\mathbf{r}_i \geq 0$, the measure $\mathbf{p} = \frac{1}{\|\mathbf{r}\|_1}\mathbf{r}$ with its approximation quality measured in $1/m_{k+1}$. As a result, to get $\|\mathbf{p}-\widehat{\mathbf{p}}\|_{\infty} \leq \delta$, it is sufficient to show:

$$\frac{6L}{M} + \frac{2\|K_{\Delta h}\|_{\infty}}{\pi\sqrt{N_z}} \leq \frac{\delta}{m_{k+1}}\|\mathbf{r}\|_1 \tag{19}$$

by choosing M and N_z such that

$$\frac{6L}{M} \leq \frac{\delta}{2m_{k+1}}\|\mathbf{r}\|_1, \qquad \frac{2\|K_{\Delta h}\|_{\infty}}{\pi\sqrt{N_z}} \leq \frac{\delta}{2m_{k+1}}\|\mathbf{r}\|_1 \tag{20}$$

assuming that for a sufficiently dense (namely, $\frac{m_k}{2} \geq m_{k-1} + \beta_k$) simplicial complex $\|\mathbf{r}\|_1 = m_k - \sum_{i=-1}^{k-1}(-1)^{k-1-i}(m_i - \beta_{i+1}) \geq \frac{1}{2}m_k$.