

Spectral Coarsening with Hodge Laplacians

Alexandros D. Keros
University of Edinburgh
Edinburgh, UK
a.d.keros@sms.ed.ac.uk

Kartic Subr
University of Edinburgh
Edinburgh, UK
k.subr@ed.ac.uk




Figure 1: We present a method to simplify geometry while preserving targeted bandwidths in a customizable spectral domain. The input to our algorithm is a *generalized spectral domain*, a combination of spectral bands from multiple discrete Laplacian operators. Our method includes a simple “lifting” operator to enable the workflow shown on the right. The example shows diffusion distances over a tetrahedral mesh of an engine part with 20 small holes that are preserved. Removing 80% of the input simplices, and seeking to preserve topology (lower portion of the spectrum), leads to some visible differences with respect to the ground truth but the quantitative comparisons indicate low error ($\approx 3.6e - 7$, shown in the supplemental material).

ABSTRACT

Many computational algorithms applied to geometry operate on discrete representations of shape. It is sometimes necessary to first simplify, or *coarsen*, representations found in modern datasets for practicable or expedited processing. The utility of a coarsening algorithm depends on both, the choice of representation as well as the specific processing algorithm or operator. e.g. simulation using the Finite Element Method, calculating Betti numbers, etc. We propose a novel method that can coarsen triangle meshes, tetrahedral meshes and simplicial complexes. Our method allows controllable preservation of salient features from the high-resolution geometry and can therefore be customized to different applications.

Salient properties are typically captured by local shape descriptors via linear differential operators – variants of *Laplacians*. Eigenvectors of their discretized matrices yield a useful *spectral domain* for geometry processing (akin to the famous Fourier spectrum which uses eigenfunctions of the derivative operator). Existing methods for spectrum-preserving coarsening use zero-dimensional discretizations of Laplacian operators (defined on vertices). We propose a generalized spectral coarsening method that considers multiple Laplacian operators defined in *different dimensionalities*

in tandem. Our simple algorithm greedily decides the order of contractions of simplices based on a quality function per simplex. The quality function quantifies the error due to removal of that simplex on a chosen band within the spectrum of the coarsened geometry.

CCS CONCEPTS

• Computing methodologies → Shape analysis.

KEYWORDS

geometry processing, numerical coarsening, spectral geometry

ACM Reference Format:

Alexandros D. Keros and Kartic Subr. 2023. Spectral Coarsening with Hodge Laplacians. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings)*, August 6–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3588432.3591544>

1 INTRODUCTION

Discrete representations of shape are ubiquitous across computer graphics applications. Meshes are specific instances of general abstractions called *simplicial complexes*. While the vertices of a mesh are commonly embedded (have explicit coordinates) in 2D or 3D, simplicial complexes capture abstract relationships between nodes – as extensions of graphs by including 3-ary (triangles), 4-ary (tetrahedra) and higher dimensional-relationships. We develop an algorithm to coarsen simplicial complexes of arbitrary dimensionality.

The choice of discretization has a profound impact on downstream applications that operate on the geometry, both in terms of efficiency as well as numerical stability. *Simplification schemes*

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 the author(s) 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.

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0159-7/23/08...\$15.00

<https://doi.org/10.1145/3588432.3591544>

are used to reduce the number of discrete elements while preserving quality. Quality can be defined either in terms of aesthetic appeal (direct rendering, visualization, etc.) or functionally (finite element simulation, topological data analysis, geometry processing algorithms, etc.). We focus on the latter and propose a coarsening algorithm that can be tailored to different applications via a versatile definition of functionally salient features.

We resort to classical spectral theory to formalize the definition of qualities of the input representation which should be preserved while coarsening. Just as the famous Fourier spectrum is obtained via projection onto the eigenfunctions of the univariate derivative operator, our spectral domain of choice is defined via a projection onto a Laplacian operator. Unlike the univariate derivative operator, a variety of definitions exist for Laplacians, which impact the utility of the coarsened mesh in applications. For example, simplification of the domain of a physics simulation requires the preservation of the spectra of *both 0- and 1-dimensional Laplacians* [De Witt et al. 2012]. The input to our algorithm includes a list of Laplacians and associated spectral bands – spectral subspaces of these Laplacians. Its output is a coarsened representation that strives to preserve the spectral profile in the specified bands.

To summarize, we propose a coarsening algorithm for simplicial complexes that can preserve spectral bands of different Laplacians, across different dimensionalities of simplices. Our simple algorithm operates by first evaluating a quality function per simplex, which quantifies the error introduced by contracting (eliminating) that simplex towards the specified spectral band(s) to be preserved. Then, we greedily perform contractions iteratively by choosing candidates with minimum error until the target coarsening level is reached. Thus, our coarsening algorithm is agnostic to the specific Laplacian considered. Our contributions in this paper are:

- a novel coarsening operator that is Laplacian-independent;
- a coarsening operator that simultaneously preserves spectral bands associated with multiple Laplacians;
- an algorithm for band-pass filtering of simplicial complexes.

We evaluate our method using a variety of surface (triangular) and volumetric (tetrahedral) meshes, as well as simplicial complexes.

2 RELATED WORK

Graphs. The spectrum of a combinatorial graph Laplacian reveals fundamental geometric and algebraic properties of the underlying graph [Chung 1999]. Several works attempt to preserve spectral subspaces, while reducing the size of input graphs [Chen et al. 2022]. A notable example [Loukas 2019] proposes an iterative, parallelizable solution to preserve spectral subspaces of graphs by minimizing undesirable projections.

Meshes. Seminal works for coarsening triangle meshes propose localized and iterative operations via edge collapses based on geometric criteria [Garland and Heckbert 1997; Ronfard and Rossignac 1996]. Similar methods have also been applied to tetrahedral mesh simplification, based on volume, quadric-based, and isosurface-preserving criteria [Chiang and Lu 2003; Chopra and Meyer 2002; Vo et al. 2007]. Recent methods formulate coarsening as an optimization problem subject to various sparsity conditions [Liu et al. 2019], by detaching the mesh from the operator [Chen et al. 2020] and

localizing error computation to form a parallelizable strategy [Lescot et al. 2020]. The *cotan Laplacian* is a popular choice, and is used, via its functional maps [Ovsjanikov et al. 2016], to identify correspondences between partial meshes [Rodolà et al. 2017].

Simplicial complexes and computational topology. The *link condition* [Dey et al. 1998] is a combinatorial criterion ensuring homology preservation while performing strong collapses (merging vertices), extending to persistent homology [Boissonnat and Pritam 2019; Wilkerson et al. 2013]. Edge collapses (edge removals) [Boissonnat and Pritam 2020; Glisse and Pritam 2022] is a state-of-the-art method for simplifying filtered simplicial complexes while preserving their (persistent) homology. These methods focus on nullspace dimensionality (kernel) of the Laplacian, rarely investigating the spectral profile of the reduced complex. Notable exceptions [Black and Maxwell 2021; Hansen and Ghrist 2019; Osting et al. 2017] apply the method of effective resistances [Spielman and Srivastava 2011] for coarsening complexes and sheaves, while Morse theory [Ebli et al. 2022] enables signal compression and reconstruction.

Spectral analysis. Laplacians make frequent appearances across geometry processing, machine learning and computational topology. Specific definitions and flavours vary widely across discrete exterior calculus [Crane et al. 2013; Desbrun et al. 2005], vector-field processing [de Goes et al. 2016; Poelke and Polthier 2016; Vaxman et al. 2016; Wardetzky 2020; Zhao et al. 2019b], fluid simulation [Liu et al. 2015], mesh segmentation and editing [Khan et al. 2020; Lai et al. 2008; Sorkine et al. 2004], topological signal processing [Barbarossa and Sardellitti 2020], random walks [Lahav and Tal 2020; Schaub et al. 2020], clustering and learning [Ebli et al. 2020; Ebli and Spreemann 2019; Keros et al. 2022; Nascimento and De Carvalho 2011; Smirnov and Solomon 2021; Su et al. 2022]. Their ability to effectively capture salient geometric, topological, and dynamic information makes their spectrum a versatile basis.

The de-facto Laplacian operator used in mesh processing is the Laplace-Beltrami operator, approximated via its discretization, the *cotan Laplacian* defined on vertices (0D), with vertex and edge weights. A multitude of other Laplacians accommodate non-manifold meshes [Sharp and Crane 2020], FEM simulations [Ayoub et al. 2020], digital surfaces [Caissard et al. 2019], polygonal meshes [Bunge et al. 2021] and arbitrary simplicial complexes [Ziegler et al. 2022]. Similarly, the spectrum of the *total variation functional* [Fumero et al. 2020] and its corresponding flow allows for spectral filtering of functions on and features of manifolds.

Motivation. The variety of definitions and properties of Laplacian operators (see Figure 1) suggests that a unified spectral coarsening algorithm could impact a range of applications. Our Laplacian-agnostic spectral coarsening can be tailored by considering the weightings as special cases of *Hodge Laplacians*. We achieve this by adapting the cost function proposed for graph theory [Loukas 2019] to simplicial complexes.

3 BACKGROUND


3.1 Simplicial Complexes & Meshes

A *simplicial complex* K is constructed from appropriate subsets of a finite set V of *vertices*. Each element $v \in V$ exists in K as a singleton

set $\{v\}$. K also contains $\sigma = \{v_0, \dots, v_k\} \subset V$ which is called a k -simplex of dimension $k = \dim \sigma = |\sigma| - 1$. e.g. edges ($k = 1$), triangles ($k = 2$), tetrahedra ($k = 3$), etc. For every such $\sigma \in K$, all of σ 's subsets $\tau \subset \sigma$ are also included in K . The dimension of the simplicial complex is the maximal dimension of its simplices. A graph $G = (V, E)$ is a 1-dimensional simplicial complex. A triangle mesh (V, E, F) is a 2-dimensional embedded simplicial complex often with additional manifold conditions on the adjacencies of 2-simplices (faces) F . A tetrahedral mesh is a 3-dimensional embedded simplicial complex.

For a graph $G = (V, E)$, the incidence matrix $A : E \rightarrow V$ where $A_{v,e} = 1$ (or -1) depending on whether $e : v \rightarrow v'$ (or $e : v' \rightarrow v$), represents directed connectivity between vertices and edges. *Boundary matrices* $\partial_k : \mathbb{R}[K_k] \rightarrow \mathbb{R}[K_{k-1}]$ extend this idea to higher dimensions and capture connectivity between (the vector space with real coefficients spanned by) the k -simplices K_k and their bounding K_{k-1} simplices. *Boundary operators* are constructed by imposing an ordering on the vertices $V = K_0$, such that each k -simplex can be expressed by an *ordered* list $\sigma = [v_0, \dots, v_k]$.

The orientation of the simplices in a simplicial complex is dictated by the ordering imposed on the vertices, and the orientation of mesh elements given by the cyclic ordering of vertices. The boundary operator acts on each simplex σ according to $\partial_k(\sigma) = \sum_{i=0}^k (-1)^i \sigma_{-i}$, where $\sigma_{-i} := [v_0, \dots, \hat{v}_i, \dots, v_k]$ indicates the deletion of the i -th vertex from σ , resulting in a $(k-1)$ -dimensional bounding simplex. We illustrate a complex along with its three boundary matrices: white cells are zeros and grey cells are ± 1 .



3.2 Hodge Laplacian as a generalization

Hodge Laplacians [Rosenberg and Steven 1997] are differential operators that extend the notion of the well-known graph and mesh Laplacians to general simplicial complexes, acting as maps from k -simplices to k -simplices. For each dimension k , they are formed as the sum of two maps: one mapping down to $(k-1)$ -simplices and another mapping up to $(k+1)$ -simplices, according to appropriate boundary matrix compositions:

$$L_k^w = L_k^{\text{down}} + L_k^{\text{up}} = \partial_k^T W_{k-1}^{-1} \partial_k W_k + W_k^{-1} \partial_{k+1} W_{k+1} \partial_{k+1}^T.$$

W_k are diagonal matrices that contain a weight per k -simplex. We direct the reader to key works [Horak and Jost 2013; Lim 2020] for an analysis of their spectral properties. We refer to L_k^w as the weighted k -Hodge Laplacian and to its unweighted (unit weights) version L_k as simply the k -Hodge Laplacian.

Most variants of Laplacians used in graph- and mesh-processing may be derived as special cases of the weighted k -Hodge Laplacian. The graph Laplacian is a 0-Hodge Laplacian with unit weight on vertices $L^{\text{graph}} = L_0^w = L_0^{\text{up}} = \partial_1 W_1 \partial_1^T$. The *cotan Laplacian* is the

0-Hodge Laplacian with weights on vertices and edges:

$$w_{v_i} = \sum_{\forall \sigma = \{v_i, v_j, v_k\} \in K_2} \mathcal{A}_\sigma / 3, \quad w_{e_{ij}} = \frac{1}{2} (\cot \theta_{lij} + \cot \theta_{mij}).$$

\mathcal{A}_σ is the area of face σ , and θ_{lij} is the angle at vertex v_l facing the edge $e_{ij} = \{v_i, v_j\}$. Figure 1 (left) depicts the impact of the choice of Laplacian on the coarsened Fertility triangle mesh.

Hodge Spectra. The basis of the homology group \mathcal{H}_k (a vector space) of a complex K is spanned by equivalence classes of k -dimensional *nontrivial loops*, and is isomorphic to the kernel of the k -Hodge Laplacian [Eckmann 1944]: $\ker L_k \simeq \mathcal{H}_k(K)$. The eigenvectors of L_k corresponding to zero eigenvalues form the harmonic part of the spectrum. An eigenvector corresponding to a non-zero eigenvalue of L_k^w , must either be an eigenvector of L_k^{up} or L_k^{down} with the same eigenvalue. Furthermore, an element of the nullspace of L_k^w must be in the kernel of both of its components. Given an eigenpair (λ_i, v_i) of L_k^{down} then $(\lambda_i, \partial_k v_i)$ is an eigenpair of L_{k-1}^{up} [Horak and Jost 2013; Torres and Bianconi 2020]. The *Hodge decomposition* ties everything together, by expressing the space of k -simplices K_k as a direct sum of *gradients*, *curls*, and *harmonics*: $K_k = \text{im } \partial_k^T \oplus \text{im } \partial_{k+1} \oplus \ker L_k$. We refer to several excellent introductions [Chen et al. 2021; Lim 2020] to Hodge decompositions.

4 METHOD

We propose a simple iterative algorithm for coarsening a simplicial complex based on two inputs: the fraction of simplices to be reduced and the portion of the input spectrum to be preserved. We use the latter to construct a quality function that is evaluated at each simplex. Then we greedily contract a simplex (or a group of simplices with low spectral-quality scores). We recalculate the quality function for simplices in the coarsened complex and iterate until the specified number of simplices have been reduced.




Figure 2: An illustrated overview of our algorithm.

4.1 Quality function

Each contraction (e.g. edge collapse) can be defined as a projection $P_k : K_k \rightarrow \hat{K}_k$ mapping k -simplices in the input complex to those in the coarsened complex. Ideally, the eigenspace to be preserved should be perpendicular to that induced by the contraction. Intuitively the projection of the former onto the latter measures “spectral leak”, with a large value indicating that the contraction leads to loss of fidelity with respect to the specified target spectrum.

Let the eigenspace to be preserved be represented by eigenvectors U and eigenvalues S (diagonal matrix) of the Hodge Laplacian L_k^w (or L_k) of the input complex, and let $A_k = U\sqrt{S^+}$ be the preserved subspace. Since P and its pseudoinverse P_k^+ map in opposite

directions to and from the output complex, the operator $\Pi = P_k^+ P_k$ projects a signal defined on the fine complex down to the coarsened complex and back up to the fine complex.

Ideally, we seek contractions where $\Pi_k A_k$ is identical to A_k by minimizing $\|\Pi_k A_k - A_k\|_\bullet$. Our quality function $c_k : K_k \rightarrow \mathbb{R}$

$$c_k = \|\Pi_k^\perp A_k\|_L, \quad (1)$$

where the perpendicular projection $\Pi_k^\perp = (\mathbb{I} - \Pi_k)$, quantifies spectral error. Here $\|x\|_L = \sqrt{x^T L x}$ denotes the L -norm. The definition of c in terms of Π_k^\perp allows for localized error computations, facilitating parallelization, since the nonzero entries of Π_k^\perp only refer to simplices affected by the contraction. This quality function has desirable theoretical properties when applied to graph simplification and general positive semidefinite matrices [Loukas 2019].

4.2 Contraction

The first step to performing contractions is to identify a set Φ of candidate contraction sets $\phi_i \in \Phi$. We execute edge collapses by setting each ϕ_i to be an edge. More complicated contractions, such as collapsing stars around vertices, may be performed by identifying the relevant $\Phi = \{\phi_i\}$.

Then, we evaluate the quality function over each candidate set and greedily contract the ϕ^* with minimum quality, identifying all the simplices in ϕ^* to one, target, vertex. As a result, the complex is coarsened while best preserving the spectral band (U, S) of L_k . We apply such contractions iteratively until a specified ratio of simplices has been contracted, with some additional bookkeeping at each level ℓ : starting from $A_k^{\ell=0} = U \sqrt{S^+}$, we update the corresponding projection and target subspace matrices P_k^ℓ and A_k^ℓ , the latter tracking the evolution of the desired spectrum at each iteration.

4.3 Multiple Laplacian subspaces

The quality function (Equation 1) can be adapted so that the target subspace is shaped via multiple spectral bands. We define cost functions $c_k^{\ell, v}$ at level ℓ independently for B different spectral bands $v = 1, 2, \dots, B$. Each of these is associated with a potentially different Laplacian. The aggregate cost function can then be tailored based on the specific downstream application as $q_{\text{agg}} : \mathbb{R}^B \rightarrow \mathbb{R}$ to obtain the final quality function $c_k^\ell = q_{\text{agg}}(\{c_k^{\ell, v}\})$, $v \in [0, 1, \dots, B]$.

In our experiments in Section 5, where L_0, L_1 and L_2 are considered in tandem, we simply average the contributions of the different Laplacians $q_{\text{agg}}(c_k^v) = \frac{1}{B} \sum_{k,v} c_k^v$. Algorithm 1 shows our generalized coarsening procedure by assembling the above stages.

4.4 Implementation details

Terms of the Hodge Laplacian. Due to the interplay of spectra of L_k and L_{k+1} discussed in Section 3.2, constraints on Laplacians for multiple k can potentially introduce ‘spectral conflicts’. We avoid this by considering only the L_k^{up} components of the chosen Laplacians for each k . For wide spectral bands and appropriate choice of k , the spectral region of interest will be a subspace of the spectrum of L_k^{up} and thus preserved.

Building coarsening matrices. Matrices P_k^\mp encode simplicial maps $\phi : K_k \rightarrow \hat{K}_k$ where every k -simplex of K maps to a valid simplex

ALGORITHM 1: Iterative generalized spectral coarsening

inputs :
 $K = \{K_k\};$ // high-resolution simplicial complex
 $\rho;$ // fraction of simplices to reduce
 $\{(S_0, U_0), \dots, (S_B, U_B)\};$ // target subspaces to preserve
 $L = \{L_0, \dots, L_B\};$ // Laplacians corresponding to subspaces

output :
 $\hat{K};$ // Coarsened complex
 $P_k;$ // projection matrices (fine to coarse)

$\hat{K} \leftarrow K;$
 $\ell \leftarrow 0;$
while $1 - |\hat{K}|/|K| < \rho$ **do**
 Define $\Phi;$ // say, set of all edges of \hat{K}
 foreach subspace (S_v, U_v) **do**
 if $\ell = 0$ **then**
 $M_v^\ell \leftarrow U_v S_v^{+1/2};$
 $A_v^\ell \leftarrow M_v^\ell;$
 else
 $L_v^\ell \leftarrow P^{\ell\mp} L_v^{\ell-1} P^{\ell+};$ // coarsened Laplacian
 $M_v^\ell \leftarrow P^\ell M_v^{\ell-1};$
 $A_v^\ell \leftarrow M^\ell (M_v^{\ell T} L_v^\ell M_v^\ell)^{+1/2};$ // coarse target subspace
 foreach set of simplices $\phi \in \Phi$ **do**
 $\Pi^\perp \leftarrow (\mathbb{I} - P^+ P);$
 foreach subspace (S_v, U_v) **do**
 $c_{\phi}^{\ell, v} \leftarrow \|\Pi^\perp A_v^\ell\|;$
 $c_\phi^\ell \leftarrow q_{\text{agg}}(\{c_{\phi}^{\ell, v}\});$ // user-specified aggreg. function
 $(\hat{K}, P^\ell) \leftarrow \text{Contract}(\hat{K}, \text{argmin}_\phi \{c_\phi^\ell\});$
 $\ell \leftarrow \ell + 1;$

of \hat{K} , or to zero. The simplices affected by the collapse of an edge $e = \{v_s, v_t\}$, where v_t is set to v_s , are the ones having e as a face, i.e. belonging to the star of the closure of e , denoted $\text{st}(\bar{e})$. Starting with the trivial projection, where P_k^\mp is an identity matrix, we modify its entries according to simplex identifications induced by the simplicial map ϕ . For each k -dimensional simplex $\sigma_k \in \text{st}(\bar{e})$ we consider two cases. If both vertices $v_s, v_t \in \sigma_k$ then the simplex is deleted and $P_{\sigma_k, \sigma_k}^\mp = 0$. If only one of the involved vertices, namely v_t , belongs to σ_k , we are faced with two alternatives, either the new simplex $\sigma'_k = \phi(\sigma_k)$ already exists in the complex K_k , or it is a newly created simplex that is first introduced in \hat{K}_k . In the former case where $\sigma'_k \in K_k$, we set $P_{\sigma'_k, \sigma_k}^\mp = 1$ and $P_{\sigma'_k, \sigma'_k}^\mp = 1$, while $P_{\sigma_k, \sigma_k}^\mp = 0$. Otherwise, if σ'_k did not previously exist in K_k , and is only introduced in \hat{K}_k , then σ_k just necessitates an appropriate reindexing of its constituent vertices. Finally, we remove zero rows from P^\mp . Theorem 4.1 (proof in the supplementary material) guarantees consistency across dimensions, and applies to contractions of larger (connected) families of simplices (not just edge collapses) since they can always be decomposed to a sequence of edge collapses.

THEOREM 4.1. *Coarsening matrices P_k^\mp commute with boundary operators*


$$\hat{\partial}_k P_k^\mp = P_{k-1}^\mp \partial_k. \quad (2)$$

For the spectral approximation guarantees [Loukas 2019] to apply to our generalized coarsening matrices, and for $P^+ = P^T D^{-2}$ to hold, we set the elements of each contraction set Φ of affected simplices as $P_{\sigma'_k, \sigma_k} = \frac{1}{|\Phi|}$, for all $\sigma_k \in \Phi$ and $\sigma'_k \in P^\mp(\Phi)$, unless it is a deleted simplex where $P_{\sigma_k, \sigma_k} = 0$. We resolve the ambiguity of the choice of target simplex by consistently selecting the target simplex as the one with minimum index.

Choosing candidate families Φ . Our algorithm is agnostic to the choice of combinations of simplices Φ to be contracted. We tested with various candidate families: edges (pairs of vertices), faces (triplets of vertices) and more general vertex neighborhoods consisting of closed-stars. Larger contraction sets result in aggressive coarsening at each iteration which leads to larger spectral error. All results in this paper use only edge collapses. This also simplifies comparison to related work (which are restricted to edge collapses).

Harmonic subspaces. For $k > 0$, the harmonic portion of the spectrum of L_k is non-trivial and encodes information about non-trivial cycles called *homology generators*. A homology group \mathcal{H}_k is a vector space of cycles that are not bounding higher dimensional simplices and therefore manifest as the null space of L_k . It is often desirable to preserve these eigenvectors, despite their corresponding zero eigenvalue, to maintain topological consistency. In practice, for all experiments we use modified eigenvalues $\tilde{S} = \mathbb{I} + S$ when dealing with spectra of higher-dimensional Laplacians.

Local Delaunay vertex position optimization. Edge contractions may produce self-intersections, particularly if a fixed vertex-placement policy is followed, say, at the midpoint of an edge. While there are methods to avoid this effectively [Sacht et al. 2013], we adopted a simple vertex positioning scheme for tetrahedral meshes to prevent intersections.



Our scheme operates in four steps: (1) identify the faces in the *link* of the edge being collapsed; (2) construct the Delaunay tetrahedralization of the vertices of these faces; (3) choose barycenters of the new tetrahedra that are on the same side of the faces in step (1); (4) the output vertex is the centroid of barycenters that pass the test in step (3), otherwise, we do not perform the collapse.

5 RESULTS

Unless otherwise specified, we use a common low-pass constraint as the default for experiments: to preserve the space spanned by the first 100 eigenvectors of the specified Laplacian(s).

5.1 Meshes

2D Baseline. To enable comparisons of L_1 for triangle meshes we extend previous work [Lescoat et al. 2020] as a baseline. Their method minimizes spectral error $E = \|PM^{-1}LF - \tilde{M}^{-1}\tilde{L}PF\|_{\tilde{M}}^2$, where P is a coarsening projection matrix, M is a mass matrix, L the Laplacian, and F is the spectrum of interest as a matrix of eigenvectors. A tilde above the respective notations denotes their coarsened versions. From their output coarsening matrix P_0 we

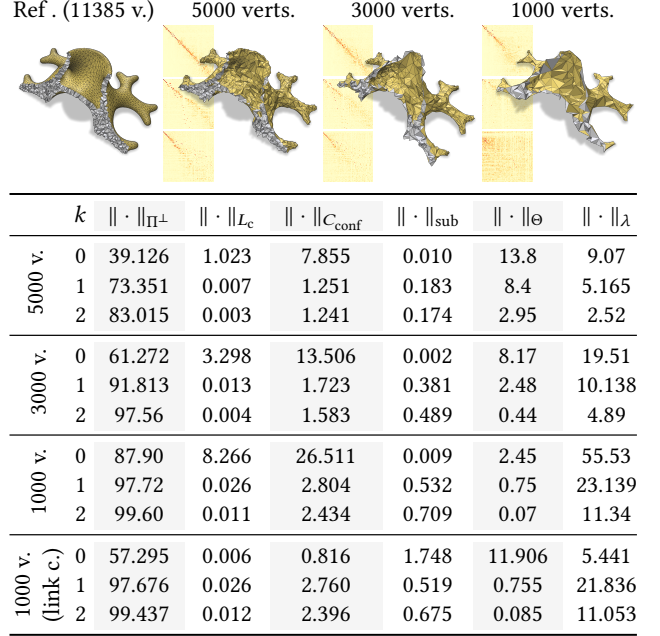


Figure 3: Spectral error increases controllably as the reference tetrahedral mesh (top left) with 11,385 vertices is progressively coarsened while preserving the spaces spanned by the first 100 eigenvectors of three Laplacians: L_k , $k = 0, 1$ and 2 .

additionally infer coarsening matrices P_1 and P_2 which operate on the space of edges and faces respectively. However, they *do not consider higher dimensional mappings in their error metric*. Since they only consider one Laplacian L_0^{cot} , we include this as one of the targets in all our comparisons, unless stated otherwise. We visualize spectral preservation via functional maps [Chen et al. 2020; Liu et al. 2019; Ovsjanikov et al. 2016] $C = U_C^T P U$ which is a diagonal matrix when the input and output spectra match perfectly. Unfortunately, it is not straightforward to compare with prior non-spectral tetrahedral coarsening methods, or to extend previous spectral coarsening methods to tetrahedral meshes (Section 2).

Quantitative comparisons. We visualize coarsened triangular (Figure 6) and tetrahedral (Figure 9) meshes along with a few eigenvectors (as heat maps on vertices). The results are reassuring that for the special case of coarsening triangle meshes our algorithm produces similar spectral results to previous work. Using L_1 subtly improves the preservation of structures such as the jaw-line, the mouth and the nose of the Suzanne model (Figure 6). The boundary loops around the eyes (eigenvectors of the null space of L_1) aim to preserve their original size compared to when only L_0^{cot} is used. Figure 6 contains quantitative comparisons of our eigenvalues against the baseline and reference. The approximation is good for L_0^{cot} , but curiously spectral divergence is observed for both methods in L_1 .

This is also observed on a tetrahedral mesh (Fertility model) as shown in Figure 9. However, the profile of eigenvalues for L_1 and the eigenvalues of L_2 are approximated well. We also show

functional maps to support these observations. The Laplacians considered were L_0^{cot} and L_1 for the triangle mesh (2D), and additionally L_2 for the tetrahedral mesh. We report errors in Figures 6 and 9 using the following error metrics: a local volume-preserving measure $\|C\|_{\Pi^\perp} = \|C^T C - \mathbb{I}\|_F^2$, an isometry measure $\|C\|_{L_c} = \|S_c C - C S\|_F^2 / \|C\|_F^2$, a conformality measure $\|C\|_{\text{conf}} = \|C^T S C - S_c\|_F^2$, a subspace approximation measure $\|C\|_{\text{sub}} = \|\|S^{1/2} U^T \Pi U S^{-1/2}\|_2 - 1\|$, subspace alignment $\|C\|_\theta = \|\sin \Theta(U, P^+ U_c)\|_F^2$ and relative eigenvalue error $\|S\|_\lambda = \|(S - S_c)/S\|_2$.

The above results indicate comparable performance against the baseline for triangle meshes and the added capability of handling higher dimensional Laplacians for tetrahedral (3D) meshes. Additional examples are included in the supplementary material. We illustrate the evolution of spectral errors for the Tree root model over three levels of coarsening (Figure 3).

5.2 Simplicial Complexes

We compare our coarsening algorithm to the state-of-the-art *edgecollapser* of the Gudhi library [The GUDHI Project 2015], that guarantees homology preservation, and a simple baseline where edges are collapsed at random.

Band-pass filtering. We demonstrate the versatility of our method by coarsening an input complex with two different spectral bands to be preserved: $\beta_1 + 1$ lowest, and 10 largest eigenpairs of L_1 . Here, β_1 is the betti number (rank of the homology group \mathcal{H}_1). Figure 4 presents these results along with the functional maps C as explained above. As expected (by design), Gudhi exactly preserves the homology rank β_1 , with no regard for other features. On the contrary, our method can be tuned to a spectral band of choice. Either the nullspace-encoded harmonic information (first row of functional maps) or the high frequency band (second row of maps).

Homology preservation. We constructed a diverse dataset of 100 simplicial complexes [Keros et al. 2022] with non-trivial topology by randomly sampling 400 points repeatedly on multi-holed tori, and subsequently constructing *alpha* complexes at various thresholds. It contains from zero up to 66 homology cycles. On this dataset, we computed mean spectral error metrics (Table 1), alongside a homology preservation error $E_\beta = |\beta_1^{\text{fine}} - \beta_1^{\text{coarse}}|$. Intuitively, this is the error in the number of 1-cycles destroyed by coarsening. The number of simplices reduced by Gudhi was consistent across both experiments. We ran two variants of our method which preserved the first 30 eigenpairs of L_1 , reducing the input complexes by a factor of 0.8, and the first $\beta_1 + 1$ eigenpairs of L_1 , matching the target number of vertices to the result of Gudhi. The results are summarized in Table 1 with standard deviations in parentheses. Gudhi, as designed, preserves β_1 exactly but exhibits spectral leak elsewhere. With random contractions, the leak is amortized across spectral bands but it destroys about 1 cycle on average. Our method is controllable, highlighting that it can be particularly effective if the spectral constraints are known for a particular application.

5.3 Applications

Denoising. A straightforward application of our method is to suppress frequencies associated with noise. We filtered a noisy Bunny model (43K vertices) using a very narrow low-pass filter:

Table 1: Spectral approximation metrics averaged over 100 complexes. Our algorithm preserves the first 30 and the first $\beta_1 + 1$ dimensions of the eigenspace respectively. The various error metrics (columns) are described in Section 5.2. Standard deviations are shown in parentheses.

k		ρ	$\ \cdot\ _{L_c}$	$\ \cdot\ _{\Pi^\perp}$	$\ \cdot\ _{\text{sub}}$	$\ \cdot\ _\lambda$	E_β
30	Gudhi	0.9	3.84 (4.2)	29.1 (2.9)	2.5 (16.8)	71.5 (664)	0.0 (0.0)
	Ours	0.8	0.49 (0.5)	8.98 (3.5)	1.52 (12.2)	2.76 (10.5)	0.07 (0.2)
	Random	0.8	3.08 (2.2)	20.8 (3.5)	0.32 (0.4)	400000 (3 e6)	.98 (1.2)
$\beta_1 + 1$	Gudhi	0.9	-	2.94 (1.0)	0.91 (0.2)	-	0.0 (0.0)
	Ours	0.9	-	1.78 (1.2)	0.78 (0.3)	-	0.21 (0.6)
	Random	0.9	-	2.76 (1.1)	0.88 (0.2)	-	1.29 (1.4)

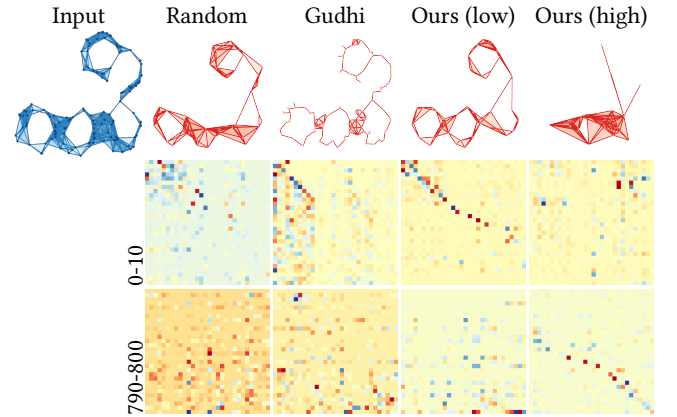


Figure 4: Low-pass and high-pass filtering of a simplicial complex using our method. Top row shows the input and coarsened complexes. Rows 2 and 3 depict the functional maps of the lowest 10 and highest 100 frequencies of L_1 . Gudhi is designed to preserve homology so its map is only diagonal for the first 4 (number of holes) elements. Our algorithm can controllably coarsen the input complex.

only the first three eigenvectors of L_0 and L_1 . Figure 7 visualizes our result along with the baseline, which filters unevenly, possibly due to noisy geometric information in its cotan weighting.

Finite Element Method (FEM). We solve the Poisson equation $-\Delta u = 1$ with Dirichlet boundary condition $u = 0$ on the Plate-hole model, using piece-wise linear (triangular) elements. In addition to solving it on the discretized mesh, we test robustness by applying planar perturbations to the vertices with increasing levels of Gaussian noise. For each setting, we coarsen the noisy mesh with $\rho = 0.75$ (8000 to 2000 vertices) using L_0^{cot} and L^1 , run a standard FEM solver on the coarse mesh and lift the solution to the input

mesh. Figure 8 shows a plot of error vs noise (top left), the reference solution (top right) and error maps for different noise settings (columns) and methods (rows). We compare our method against previous work [Lescoat et al. 2020], with (optpos) and without (noopt) vertex position optimization, a quadric-based method [Garland and Heckbert 1997] and a uniform mesh decimation baseline where edges are collapsed uniformly at random (Uniform). We further compare against a multigrid approach [Liu et al. 2021] via its full solver (MG-F) versus just the prolongation operator for lifting our computed Poisson solution (MG-P) to the original mesh. The mesh is coarsened with QSLim [Garland and Heckbert 1997], a shortest-edge midpoint collapse strategy (MP), and a vertex removal strategy (VR). The coarse mesh is overlaid on the error maps and can be viewed by magnifying the figure. Our method exhibits robustness against noise when compared to the baselines. Although the specialized (multigrid) solver [Liu et al. 2021] performs better at low noise levels, it fails to reach the desired number of vertices when the standard deviation of Gaussian noise is increased beyond 0.5.

Spectral distances. We evaluate the fidelity (Figure 10) of spectral distance measures computed on the coarse mesh and “lifted” to the fine mesh, between vertices w, v of a triangle surface mesh (Suzanne) and a volume tetrahedral mesh (Fertility), using the same parameters as Figures 6 and 9. Figure 1 illustrates the spectral distance approximation on the *Engine* tetrahedral model, coarsened with $\rho = .8$ (from 46220 to 10000 vertices) while preserving the first 50 eigenvectors of L_0^{cot} and the first 25 eigenvectors of L_1 and L_2 . The model has 20 holes that are retained in the coarse mesh. The metrics used are the diffusion distance, the biharmonic distance, the commute distance, the wave kernel signature and distance [Aubry et al. 2011], and the heat kernel signature and distance [Sun et al. 2009]. Formulas are provided in the supplementary material.

Spectral distance approximation on Suzanne outperforms the baseline by orders of magnitude in most cases, with error remaining low even for the tetrahedral Fertility mesh. The qualitative comparison indicates agreement between our “lifted” version and the reference, despite some localized distortions.

6 DISCUSSION

Execution time. We implemented our method in C++, and performed experiments on a 16-core workstation (Intel E5-2630 v3, 2.4 GHz) with 64GB RAM. Our method compares favorably (orders of magnitude faster) against the spectral coarsening baseline in Figure 5. Although it appears that we are competitive with the quadric-based method [Garland and Heckbert 1997]) it should be noted that ours (like the baseline) requires eigenspaces of Laplacians as input while the latter does not. The time for the eigendecomposition is not reflected in the plot (which only measures geometric operations) and should be added to both our method and the baseline. The triangle meshes used in the comparison contain 10772 (Fertility) and 29690 (Dinoskull) vertices, respectively. We use L_0^{cot} for the baseline, and L_0^{cot} & L_1 for our method.

Boundary. The intricacies of boundary values problems involving Hodge Laplacians [Mitrea 2016] complicate practical computations. The differential operators need to be constructed carefully to guarantee coverage of tangential and normal conditions [Zhao




Figure 5: Computation time vs reduction ratio.

et al. 2019a]. Since our method is Laplacian-agnostic, any positive semi-definite operator, with its spectrum, can be provided as input.

Implementation of vertex positioning. We used the *Tetgen* software [Hang 2015] for local tetrahedralization. We observed stray cases where Tetgen would not terminate, due to degeneracies. We jittered the vertices of the link of the collapsing simplex by a small ($1e-5$) amount to resolve this. Despite this, extreme simplification can cause occasional (0.01% of initial tetrahedra) intersections (at $1e-5$ detection tolerance), which we fix via extra collapses.

Link condition. We do not explicitly enforce the *link condition* [Dey et al. 1998] in our experiments, which explains the non-zero error in experiments on homology preservation. However, we found that explicit enforcement does not necessarily limit the extent to which coarsening can be performed, possibly due to the greedy nature of our algorithm, e.g. the last set of results in Figure 3 show errors under coarsening with $\rho = 0.9$, and the link condition enforced.

Lifting operator. Although our method includes a simple lifting operator, we do not expect it to outperform sophisticated operators that construct functional correspondences between meshes [Jiang et al. 2020; Liu et al. 2021; Panozzo et al. 2013]. However, our operator exhibits two desirable properties: It generalizes vertex correspondence maps to arbitrary dimensions and simplex-spanned vector spaces, and it does not involve numerical pseudoinversion, since it is assembled during spectral coarsening, which is our primary focus. Projection [Jiang et al. 2020] and prolongation [Liu et al. 2021] operations between fine and coarse geometry depend on the decimation strategy. It is interesting to investigate the interplay between such approaches and our method in future work.

Limitation and future work. The utility of our tool hinges on knowledge of spectral constraints imposed by downstream tasks. While this is known in some cases (homology, spectral distances and denoising), such constraints are not typically known across applications. The need for the use of higher dimensional Laplacian operators has been identified in some applications such as for PDE discretization in higher dimensions [Arnold 2018] and signal analysis over arbitrary simplicial complexes [Barbarossa and Sardellitti 2020]. However, we are hopeful that the availability of a tool such as ours will inspire the graphics communities (particularly geometry processing and simulation) to explore the applicability of mixed-dimensional Laplacian spectral constraints.

7 CONCLUSION

We presented a simple and efficient algorithm for coarsening simplicial complexes, while preserving targeted spectral subspaces across multiple dimensions. We exemplified the impact of the choices of spectral domain on applications such as denoising, FEM and approximate distance calculations on coarsened meshes. We hope that this work will pave the way towards unleashing the potential of using mixtures of Laplacians in discrete geometry processing.

ACKNOWLEDGMENTS

Kartic Subr was supported by a Royal Society University Research Fellowship.

REFERENCES

- Douglas N Arnold. 2018. *Finite element exterior calculus*. SIAM.
- Mathieu Aubry, Ulrich Schlickewel, and Daniel Cremers. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*. IEEE, 1626–1633.
- Rama Ayoub, Aziz Hamdouni, and Dina Razafindralandy. 2020. A new Hodge operator in Discrete Exterior Calculus. Application to fluid mechanics. *arXiv preprint arXiv:2006.16930* (2020).
- Sergio Barbarossa and Stefania Sardellitti. 2020. Topological signal processing over simplicial complexes. *IEEE Transactions on Signal Processing* 68 (2020), 2992–3007.
- Mitchell Black and William Maxwell. 2021. Effective Resistance and Capacitance in Simplicial Complexes and a Quantum Algorithm. In *32nd International Symposium on Algorithms and Computation (ISAAC 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 212)*, Hee-Kap Ahn and Kunihiko Sadakane (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 31:1–31:27. <https://doi.org/10.4230/LIPIcs.ISAAC.2021.31>
- Jean-Daniel Boissonnat and Siddharth Pritam. 2019. Computing Persistent Homology of Flag Complexes via Strong Collapses. In *SoCG 2019-International Symposium on Computational geometry*.
- Jean-Daniel Boissonnat and Siddharth Pritam. 2020. Edge collapse and persistence of flag complexes. In *SoCG 2020-36th International Symposium on Computational Geometry*.
- A. Bunge, M. Botsch, and M. Alexa. 2021. The Diamond Laplace for Polygonal and Polyhedral Meshes. *Computer Graphics Forum* 40, 5 (Aug. 2021), 217–230. <https://doi.org/10.1111/cgf.14369>
- Thomas Caissard, David Coeurjolly, Jacques-Olivier Lachaud, and Tristan Roussillon. 2019. Laplace–beltrami operator on digital surfaces. *Journal of Mathematical Imaging and Vision* 61, 3 (2019), 359–379.
- Honglin Chen, Hsueh-Ti Derek Liu, Alec Jacobson, and David I. W. Levin. 2020. Chordal Decomposition for Spectral Coarsening. 39, 6, Article 265 (nov 2020), 16 pages. <https://doi.org/10.1145/3414685.3417789>
- Jie Chen, Yousef Saad, and Zechen Zhang. 2022. Graph coarsening: from scientific computing to machine learning. *SeMA Journal* (2022), 1–37.
- Yu-Chia Chen, 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).
- Yi-Jen Chiang and Xiang Lu. 2003. Progressive Simplification of Tetrahedral Meshes Preserving All Isosurface Topologies. *Computer Graphics Forum* 22, 3 (Sept. 2003), 493–504. <https://doi.org/10.1111/1467-8659.00697>
- Prashant Chopra and Joerg Meyer. 2002. Tetfusion: an algorithm for rapid tetrahedral mesh simplification. In *IEEE Visualization, 2002. VIS 2002*. IEEE, 133–140.
- F R K Chung. 1999. Spectral Graph Theory. *ACM SIGACT News* 30 (1999), 14. <https://doi.org/10.1145/568547.568553>
- Keenan Crane, Fernando De Goes, Mathieu Desbrun, and Peter Schröder. 2013. Digital geometry processing with discrete exterior calculus. In *ACM SIGGRAPH 2013 Courses*. 1–126.
- Fernando de Goes, Mathieu Desbrun, and Yiyi Tong. 2016. Vector Field Processing on Triangle Meshes. In *ACM SIGGRAPH 2016 Courses (Anaheim, California) (SIGGRAPH '16)*. Association for Computing Machinery, New York, NY, USA, Article 27, 49 pages. <https://doi.org/10.1145/2897826.2927303>
- Tyler De Witt, Christian Lessig, and Eugene Fiume. 2012. Fluid Simulation Using Laplacian Eigenfunctions. *ACM Trans. Graph.* 31, 1 (Jan. 2012), 1–11. <https://doi.org/10.1145/2077341.2077351>
- Mathieu Desbrun, Anil N Hirani, Melvin Leok, and Jerrold E Marsden. 2005. Discrete exterior calculus. *arXiv preprint math/0508341* (2005).
- Tamal K Dey, Herbert Edelsbrunner, Sumanta Guha, and Dmitry V Nekhayev. 1998. Topology preserving edge contraction. In *Publ. Inst. Math.(Beograd)(NS. Citeseer*.
- Stefania Ebli, Michaël Defferrard, and Gard Spreemann. 2020. Simplicial neural networks. *arXiv preprint arXiv:2010.03633* (2020).
- Stefania Ebli, Celia Hacker, and Kelly Maggs. 2022. Morse Theoretic Signal Compression and Reconstruction on Chain Complexes. *arXiv preprint arXiv:2203.08571* (2022).
- 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.
- Beno Eckmann. 1944. Harmonische funktionen und randwertaufgaben in einem komplex. *Commentarii Mathematici Helvetici* 17, 1 (1944), 240–255.
- Marco Fumero, Michael Möller, and Emanuele Rodolà. 2020. Nonlinear spectral geometry processing via the tv transform. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
- Michael Garland and Paul S Heckbert. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 209–216.
- Marc Glisse and Siddharth Pritam. 2022. Swap, Shift and Trim to Edge Collapse a Filtration. *arXiv preprint arXiv:2203.07022* (2022).
- Si Hang. 2015. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw* 41, 2 (2015), 11.
- Jakob Hansen and Robert Ghrist. 2019. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology* 3, 4 (2019), 315–358.
- Danijela Horak and Jürgen Jost. 2013. Spectra of combinatorial Laplace operators on simplicial complexes. *Advances in Mathematics* 244 (2013), 303–336.
- Zhongshi Jiang, Teseo Schneider, Denis Zorin, and Daniele Panozzo. 2020. Bijective projection in a shell. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–18.
- Alexandros Keros, Vidit Nanda, and Kartic Subr. 2022. Dist2Cycle: A Simplicial Neural Network for Homology Localization. In *36th AAAI Conference on Artificial Intelligence*.
- Dawar Khan, Alexander Plopski, Yuichiro Fujimoto, Masayuki Kanbara, Gul Jabeen, Yongjie Jessica Zhang, Xiaopeng Zhang, and Hirokazu Kato. 2020. Surface remeshing: A systematic literature review of methods and research directions. *IEEE transactions on visualization and computer graphics* 28, 3 (2020), 1680–1713.
- Alon Lahav and Ayellet Tal. 2020. Meshwalker: Deep mesh understanding by random walks. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–13.
- Yu-Kun Lai, Shi-Min Hu, Ralph R Martin, and Paul L Rosin. 2008. Fast mesh segmentation using random walks. In *Proceedings of the 2008 ACM symposium on Solid and physical modeling*. 183–191.
- Thibault Lescot, Hsueh-Ti Derek Liu, Jean-Marc Thiery, Alec Jacobson, Tamy Boubekeur, and Maks Ovsjanikov. 2020. Spectral mesh simplification. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 315–324.
- Lek-Heng Lim. 2020. Hodge Laplacians on graphs. *Siam Review* 62, 3 (2020), 685–715.
- Beibei Liu, Gemma Mason, Julian Hodgson, Yiyi Tong, and Mathieu Desbrun. 2015. Model-reduced variational fluid simulation. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–12.
- Hsueh-Ti Derek Liu, Alec Jacobson, and Maks Ovsjanikov. 2019. Spectral Coarsening of Geometric Operators. *ACM Trans. Graph.* 38, 4, Article 105 (jul 2019), 13 pages. <https://doi.org/10.1145/3306346.3322953>
- Hsueh-Ti Derek Liu, Jiayi Eris Zhang, Mirela Ben-Chen, and Alec Jacobson. 2021. Surface multigrid via intrinsic prolongation. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Andreas Loukas. 2019. Graph Reduction with Spectral and Cut Guarantees. *J. Mach. Learn. Res.* 20, 116 (2019), 1–42.
- Dorina Mitrea. 2016. *The Hodge-Laplacian: Boundary Value Problems on Riemannian Manifolds*. Number volume 64 in De Gruyter Studies in Mathematics. De Gruyter, Berlin ; Boston.
- Maria CV Nascimento and Andre CPLF De Carvalho. 2011. Spectral methods for graph clustering—a survey. *European Journal of Operational Research* 211, 2 (2011), 221–231.
- Braxton Osting, Sourabh Palande, and Bei Wang. 2017. Towards spectral sparsification of simplicial complexes based on generalized effective resistance. *arXiv preprint arXiv:1708.08436* (2017).
- Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. 2016. Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*. 1–60.
- Daniele Panozzo, Ilya Baran, Olga Diamanti, and Olga Sorkine-Hornung. 2013. Weighted averages on surfaces. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–12.
- Konstantin Poelke and Konrad Polthier. 2016. Boundary-aware Hodge decompositions for piecewise constant vector fields. *Computer-Aided Design* 78 (2016), 126–136.
- Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. 2017. Partial functional correspondence. In *Computer graphics forum*, Vol. 36. Wiley Online Library, 222–236.
- Rémi Ronfard and Jarek Rossignac. 1996. Full-range approximation of triangulated polyhedra. In *Computer Graphics Forum*, Vol. 15. Wiley Online Library, 67–76.
- Steven Rosenberg and Rosenberg Steven. 1997. *The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds*. Number 31. Cambridge University Press.
- Leonardo Sacht, Alec Jacobson, Daniele Panozzo, Christian Schüller, and Olga Sorkine-Hornung. 2013. Consistent volumetric discretizations inside self-intersecting surfaces. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 147–156.

- 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.
- Nicholas Sharp and Keenan Crane. 2020. A laplacian for nonmanifold triangle meshes. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 69–80.
- Dmitriy Smirnov and Justin Solomon. 2021. HodgeNet: learning spectral geometry on triangle meshes. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–11.
- Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. 2004. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. 175–184.
- Daniel A Spielman and Nikhil Srivastava. 2011. Graph sparsification by effective resistances. *SIAM J. Comput.* 40, 6 (2011), 1913–1926.
- Xing Su, Shan Xue, Fanzhen Liu, Jia Wu, Jian Yang, Chuan Zhou, Wenbin Hu, Cecile Paris, Surya Nepal, Di Jin, Quan Z. Sheng, and Philip S. Yu. 2022. A Comprehensive Survey on Community Detection With Deep Learning. *IEEE Transactions on Neural Networks and Learning Systems* (2022), 1–21. <https://doi.org/10.1109/TNNLS.2021.3137396>
- Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. 2009. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, Vol. 28. Wiley Online Library, 1383–1392.
- The GUDHI Project. 2015. *GUDHI User and Reference Manual*. GUDHI Editorial Board. <http://gudhi.gforge.inria.fr/doc/latest/>
- Joaquin J Torres and Ginestra Bianconi. 2020. Simplicial complexes: higher-order spectral dimension and dynamics. *Journal of Physics: Complexity* 1, 1 (2020), 015002.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum* 35, 2 (2016), 545–572. <https://doi.org/10.1111/cgf.12864> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12864>
- Huy T. Vo, Steven P. Callahan, Peter Lindstrom, Valerio Pascucci, and Claudio T. Silva. 2007. Streaming Simplification of Tetrahedral Meshes. *IEEE Transactions on Visualization and Computer Graphics* 13, 1 (2007), 145–155. <https://doi.org/10.1109/TVCG.2007.21>
- Max Wardetzky. 2020. Discrete Laplace operators. *An Excursion Through Discrete Differential Geometry: AMS Short Course, Discrete Differential Geometry, January 8-9, 2018, San Diego, California* 76 (2020), 1.
- Adam C Wilkerson, Terrence J Moore, Ananthram Swami, and Hamid Krim. 2013. Simplifying the homology of networks via strong collapses. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 5258–5262.
- Rundong Zhao, Mathieu Desbrun, Guo-Wei Wei, and Yiyang Tong. 2019a. 3D Hodge Decompositions of Edge- and Face-Based Vector Fields. *ACM Trans. Graph.* 38, 6 (Dec. 2019), 1–13. <https://doi.org/10.1145/3355089.3356546>
- Rundong Zhao, Mathieu Desbrun, Guo-Wei Wei, and Yiyang Tong. 2019b. 3D Hodge decompositions of edge-and face-based vector fields. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.
- Cameron Ziegler, Per Sebastian Skardal, Haimonti Dutta, and Dane Taylor. 2022. Balanced Hodge Laplacians optimize consensus dynamics over simplicial complexes. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 32, 2 (2022), 023128.

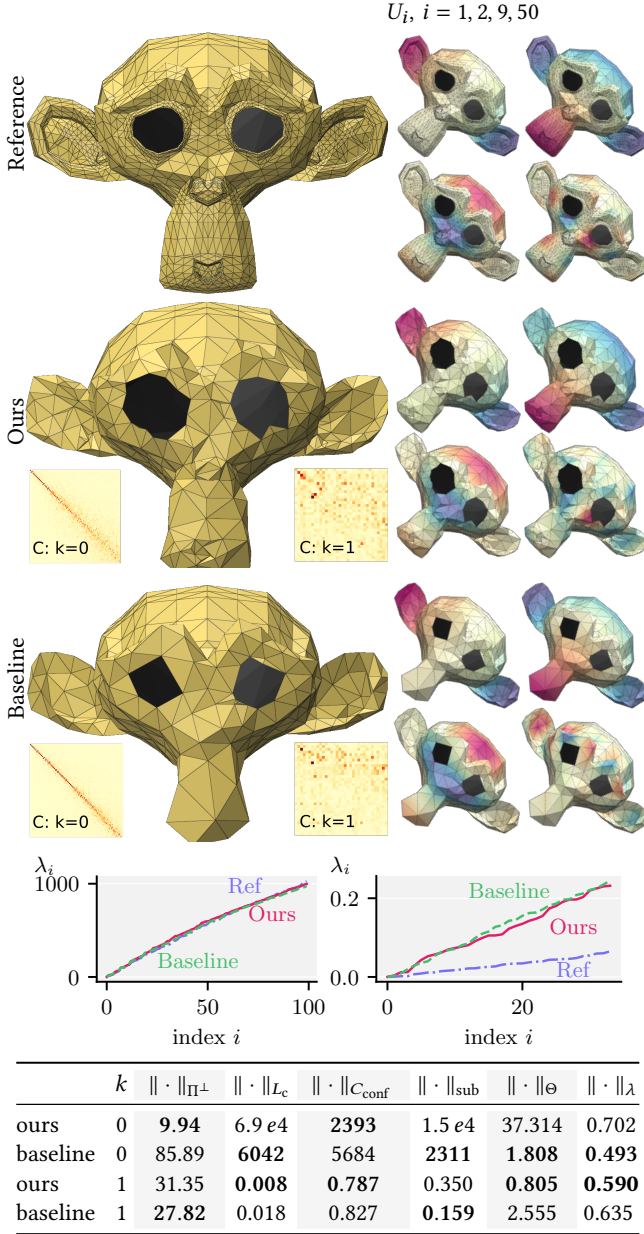


Figure 6: The input mesh, coarsened meshes (first column) and four eigenvectors of their L_0^{cot} (right) computed using our method (second row) and the baseline (third row). The mesh was coarsened from 1640 to 400 vertices, preserving L_0^{cot} and L_1 . Our method is mostly similar with subtle differences around the square jaw and button nose. The eigenvalues of the preserved spectrum are plotted below with those corresponding to $k = 0$ on the left and $k = 1$ on the right, along with spectral errors at the bottom.




Figure 7: Denoising. A noisy *Bunny* model (43645 vertices) is coarsened to 2% of its size (1000 vertices), to preserve only the first 3 eigenvectors of L_0 & L_1 using our method, and the first 3 eigenvectors of L_0^{cot} for the baseline.

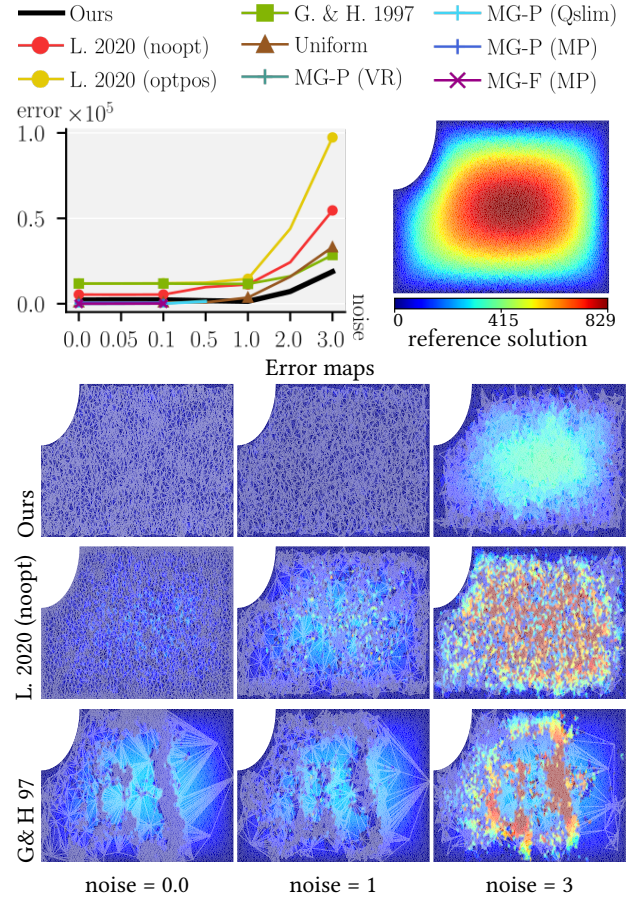


Figure 8: FEM on our coarsened mesh is robust to Gaussian noise added to the vertices (top left). The solution to $-\Delta u = 1$ on the fine mesh (zoom in to view) is shown on the top right. The table below compares error as heat maps of simulation error (blue is low and red is high), with the coarse mesh (zoom to view) overlaid, for ours and two other coarsening methods (rows).

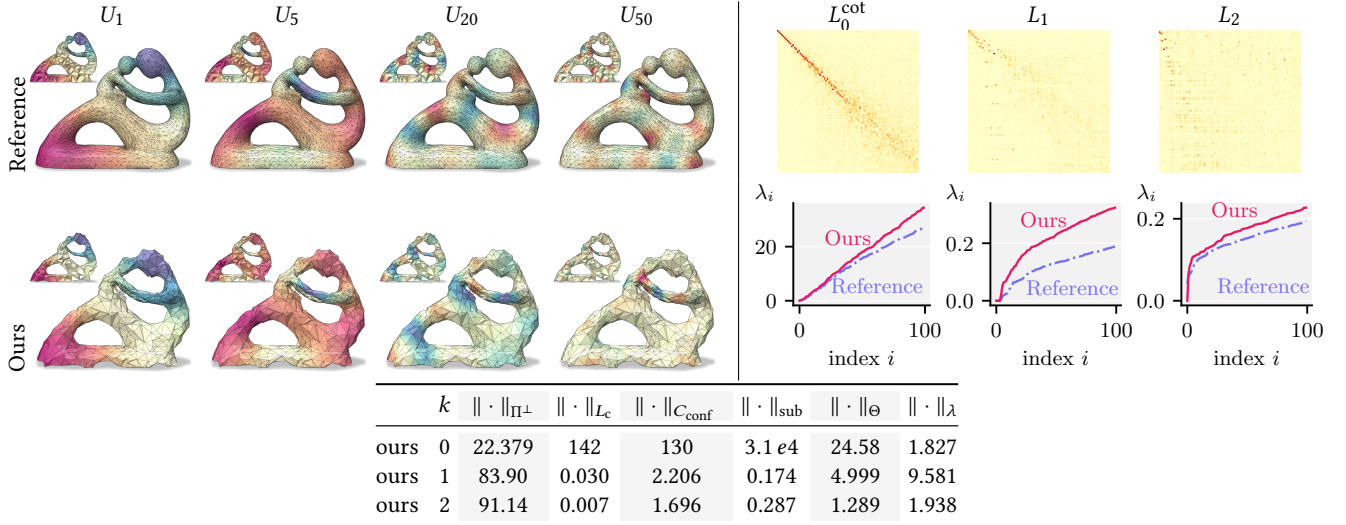


Figure 9: Four eigenvectors U_i (right) of L_0^{cot} on the fertility model coarsened from 10K to 3K vertices, preserving the first 100 eigenvectors of L_0^{cot}, L_0 and L_1 . Insets depict mesh slices, and spectral preservation is shown on the right and bottom.

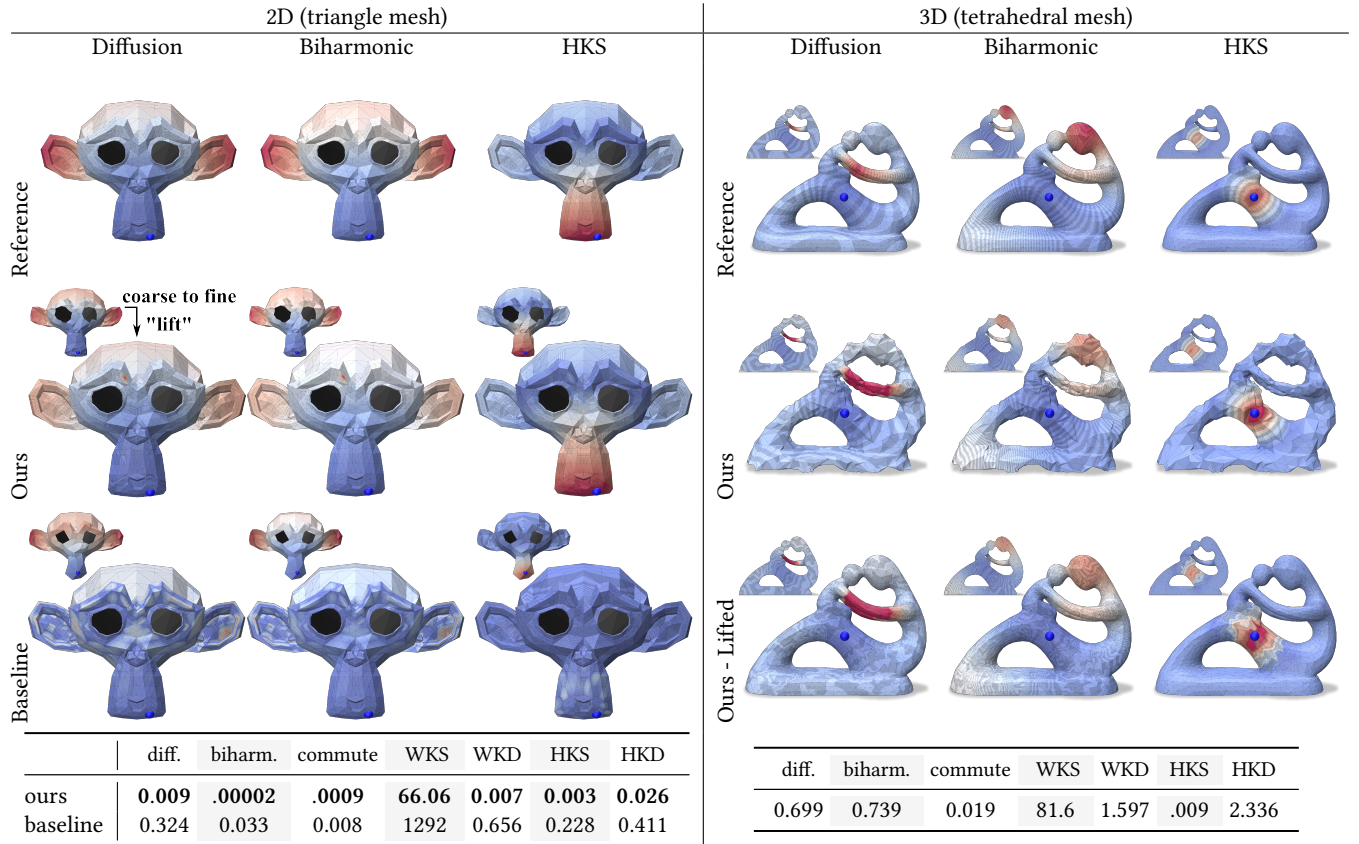


Figure 10: Different distance metrics from a point (blue dot on each mesh) are computed on a coarse mesh and visualized using colored isocontours. Our P^+ operator enables the “lifting” of these distances from the coarse mesh (insets in 2D plots and third row of 3D plots). The tables report mean squared error error for various distance metrics. The insets in the 3D plots depict slices through the mesh.