

Hodge Decomposition and General Laplacian Solvers for Embedded Simplicial Complexes

Mitchell Black ✉

School of Electrical Engineering and Computer Science,
Oregon State University, Corvallis, OR, USA

Amir Nayyeri ✉

School of Electrical Engineering and Computer Science,
Oregon State University, Corvallis, OR, USA

Abstract

We describe a nearly-linear time algorithm to solve the linear system $L_1x = b$ parameterized by the first Betti number of the complex, where L_1 is the 1-Laplacian of a simplicial complex K that is a subcomplex of a collapsible complex X linearly embedded in \mathbb{R}^3 . Our algorithm generalizes the work of Black et al. [SODA2022] that solved the same problem but required that K have trivial first homology. Our algorithm works for complexes K with arbitrary first homology with running time that is nearly-linear with respect to the size of the complex and polynomial with respect to the first Betti number. The key to our solver is a new algorithm for computing the Hodge decomposition of 1-chains of K in nearly-linear time. Additionally, our algorithm implies a nearly quadratic solver and nearly quadratic Hodge decomposition for the 1-Laplacian of any simplicial complex K embedded in \mathbb{R}^3 , as K can always be expanded to a collapsible embedded complex of quadratic complexity.

2012 ACM Subject Classification Theory of computation → Computational geometry; Mathematics of computing → Algebraic topology; Theory of computation → Design and analysis of algorithms; Mathematics of computing → Numerical analysis

Keywords and phrases Computational Topology, Laplacian solvers, Combinatorial Laplacian, Hodge decomposition, Parameterized Complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.23

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2205.02134>

Funding The authors were supported in part by NSF grants CCF-1941086 and CCF-1816442.

Acknowledgements The authors would like to thank the reviewers for their helpful comments, especially for an observation that improved the dependence on β in the runtime.

1 Introduction

The d th combinatorial Laplacian of a simplicial complex K is a linear operator that acts on vectors of real numbers associated to the d -simplices of K . The d th combinatorial Laplacian is defined as

$$L_d = \partial_d^T \partial_d + \partial_{d+1} \partial_{d+1}^T,$$

where $\partial_d : C_d(K) \rightarrow C_{d-1}(K)$ is the d th boundary map of K , and $C_d(K)$ is the d th chain group of K . The d th Laplacian encodes the incidence of $(d-1)$ -, d - and $(d+1)$ -simplices. In particular, the 0th Laplacian L_0 is composed of a constant map $\partial_0^T \partial_0$, and the well-known graph Laplacian $\partial_1 \partial_1^T$. The graph Laplacian matrix and its algebraic properties have been extensively studied in algebraic and spectral graph theory, a topic that has flourished into a rich field with many applications in computer science such as graph clustering [25, 27], graph sparsification [29], and max flow solvers [9] (see Spielman’s book and references therein [28]).



© Mitchell Black and Amir Nayyeri;

licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 23; pp. 23:1–23:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



A highlight of recent advances in algorithmic spectral graph theory is nearly-linear time solvers for linear systems on the graph Laplacian that emerged as a result of decades of research [2, 4, 11, 19, 21–23, 30, 34]. These results imply nearly-linear time solvers for the more general class of symmetric diagonally dominant matrices. They also have triggered research to find out which classes of linear systems admit nearly-linear time solvers [24]. Moreover, these solvers are used for different application areas such as approximation algorithm design and numerical analysis [5, 9].

Recent work has attempted to extend the success of graph Laplacian solvers to higher dimensional Laplacians. Cohen et al. initiated this line of work by introducing a nearly-linear solver for the 1-Laplacian of collapsible complexes embedded in \mathbb{R}^3 [10]. Black et al. continued this work by considering complexes with trivial first homology that were subcomplexes of collapsible complexes embedded in \mathbb{R}^3 [3]. The solver of Black et al. implies a nearly quadratic solver for any complex with trivial first homology embedded in \mathbb{R}^3 ; they show that a complex embedded in \mathbb{R}^3 can be extended to a collapsible embedded complex with at most quadratic complexity.

In this paper, we extend the work of Cohen et al. and Black et al. to any subcomplex of a collapsible complex embedded in \mathbb{R}^3 , regardless of the rank of its first homology group. The running time of our solver is nearly-linear with respect to the size of the collapsible complex, and polynomial with respect to the rank of its first homology group. The main tool in our paper is a new algorithm for computing the Hodge Decomposition of a 1-chain.

Computing the Hodge decomposition is a problem of independent interest since the Hodge decomposition has found a myriad of applications in topological data analysis, numerical analysis, and computer graphics among other areas [1, 8, 12, 13, 16, 20, 32, 33, 35]. The Hodge decomposition can be computed exactly in $O(n^\omega)$ time by solving a constant number of systems of linear equations, where ω is the matrix multiplication constant. (Approximately) computing the Hodge decomposition in nearly-linear time has been an open question with many possible applications.

Cohen et al. describe nearly-linear projection operators into the coboundary space and cycle space, which implies Hodge decomposition for complexes with trivial homology as the boundary and cycle spaces are identical in this case. In this paper, we describe projection operators into the boundary and harmonic spaces for an arbitrary subcomplex of a collapsible simplicial complex embedded in \mathbb{R}^3 . Our boundary projection operator is key to our solver. Our results imply 1-Laplacian solvers and projection operators for any simplicial complex embedded in \mathbb{R}^3 that are quadratic in the size of the complex and polynomial in the first Betti number; these follow from the fact that any complex in \mathbb{R}^3 can be extended to a collapsible complex in \mathbb{R}^3 with a quadratic number of simplices [3, Corollary 3.3].

While this paper presents a positive result on extending graph Laplacian solvers to a more general class of Laplacians, a recent work by Ding et al. [15] shows that solving linear equations in arbitrary 1-Laplacians (and therefore arbitrary d -Laplacians) is as hard as solving arbitrary sparse linear equations with bounded integer entries and bounded condition number. An interesting open question is whether or not there exist fast solvers for other classes of simplicial complexes.

1.1 Our Results

Let X be a collapsible simplicial complex with a known collapsing sequence embedded in \mathbb{R}^3 , and let K be a subcomplex of X . The first result of this paper is a 1-Laplacian solver for K . Recall that $L_1 = \partial_2 \partial_2^T + \partial_1^T \partial_1$. We define $L_1^{up} = \partial_2 \partial_2^T$ and $L_1^{down} = \partial_1^T \partial_1$. We refer to L_1^{up} and L_1^{down} as the up-Laplacian and down-Laplacian, respectively.

► **Theorem 1.** *Let X be a collapsible simplicial complex with a known collapsing sequence linearly embedded in \mathbb{R}^3 , and let $K \subset X$ be a subcomplex of X . For any $\varepsilon > 0$, there is an operator $\text{Laplacian.Solver}(X, K, \varepsilon)$ such that*

$$(1 - \varepsilon)(L_1[K])^+ \preceq \text{Laplacian.Solver}(X, K, \varepsilon) \preceq (L_1[K])^+.$$

where $(L_1[K])^+$ is the pseudoinverse of the 1-Laplacian $L_1[K]$. Further, for any $x \in C_1$, $\text{Laplacian.Solver}(X, K, \varepsilon) \cdot x$ can be computed in $\tilde{O}(\beta^3 \cdot n \cdot \log n \cdot \log(n/(\lambda_{\min}(K) \cdot \lambda_{\min}(X) \cdot \varepsilon)))^1$ time, where n is the total number of simplices in X , $\lambda_{\min}(K)$ and $\lambda_{\min}(X)$ are the smallest nonzero eigenvalues of $L_1^{up}(K)$ and $L_1^{up}(X)$ respectively, and β is the rank of the first homology group of K .

This result is a generalization of Theorem 1.1 of Black et al. [3] that requires K to have trivial first homology. Their running time depends on $\log(n\kappa/\varepsilon)$, with κ being the condition number of $L_1^{up}(K)$ within the boundary space. The condition number is defined $\kappa = \lambda_{\max}(K)/\lambda_{\min}(K)$, where $\lambda_{\max}(K)$ is the largest eigenvalue of $L_1^{up}(K)$, and $\lambda_{\min}(K)$ is the smallest nonzero eigenvalue of $L_1^{up}(K)$. We observe that $\lambda_{\max}(K)$ is polynomially bounded with respect to the size of the complex (Lemma 29 in the full paper.) Therefore, the log dependence of the running time of Black et al.'s solver can be simplified to $\log(n/(\lambda_{\min}(K) \cdot \varepsilon))$. The running time of Theorem 1, in contrast, has an extra dependence to $\lambda_{\min}(X)$ within the log, in addition to a polynomial dependence to β . For the special case that $\beta = 0$, we can eliminate the dependence on $\lambda_{\min}(X)$ with a more careful analysis and match the running time of Black et al.

The new ingredient that makes Theorem 1 possible is an approximate projection operator onto the boundary space. Lacking this operator, the previous papers had to assume that K has trivial homology and use a projection into the cycle space instead.

► **Lemma 2.** *Let K be a simplicial complex linearly embedded in a collapsible complex X with a known collapsing sequence that is embedded in \mathbb{R}^3 , and let Π_{bd} be the orthogonal projection operator into the space of boundary 1-chains in K . For any $\varepsilon > 0$, there is an operator $\tilde{\Pi}_{bd}(\varepsilon)$, such that*

$$(1 - \varepsilon)\Pi_{bd} \preceq \tilde{\Pi}_{bd}(\varepsilon) \preceq (1 + \varepsilon)\Pi_{bd}.$$

Further, for any 1-chain x , $\tilde{\Pi}_{bd}(\varepsilon) \cdot x$ can be computed in $\tilde{O}(\beta^3 \cdot n \cdot \log n \cdot \log(\frac{n}{\lambda_{\min}(X) \cdot \varepsilon}))$ time, where β is the rank of the first homology group of K , n is the total number of simplices in X , and $\lambda_{\min}(X)$ is the smallest nonzero eigenvalue of $L_1^{up}(X)$.

A key technical challenge to achieve our projection operator onto the boundary space is computing a projection into the space of harmonic chains, formalized in part (ii) of the following lemma. Note that our approximation guarantee for projection into the harmonic space is weaker than the one for projection into the boundary space (more on this in the overview).

► **Lemma 3.** *Let K be a subcomplex of a collapsible simplicial complex X with a known collapsing sequence that is linearly embedded in \mathbb{R}^3 . Let β be the rank of the first homology group of K , n be the total number of simplices in X , and $\lambda_{\min}(X)$ be the smallest nonzero eigenvalue of $L_1^{up}(X)$.*

¹ The $\tilde{O}(\cdot)$ notations hides a factor of $\log \log n$.

- (i) For any $\varepsilon > 0$, there is an $\tilde{O}(\beta^2 \cdot n \cdot \log n \cdot \log(\frac{n}{\lambda_{\min}(X) \cdot \varepsilon}))$ time algorithm to compute an orthonormal set of vectors $\{\tilde{g}_1, \dots, \tilde{g}_\beta\}$ such that there exists an orthonormal harmonic basis $\{g_1, \dots, g_\beta\}$ with $\|g_i - \tilde{g}_i\| \leq \varepsilon$ for all $1 \leq i \leq \beta$.
- (ii) For any $\varepsilon > 0$, there exists a symmetric matrix $\tilde{\Pi}_{hr}(\varepsilon)$ such that,

$$\Pi_{hr} - \varepsilon I \preceq \tilde{\Pi}_{hr}(\varepsilon) \preceq \Pi_{hr} + \varepsilon I,$$

where Π_{hr} is the orthogonal projection into the harmonic space. Moreover, for any 1-chain x , $\tilde{\Pi}_{hr}(\varepsilon) \cdot x$ can be computed in $\tilde{O}(\beta^2 \cdot n \cdot \log n \cdot \log(\frac{n}{\lambda_{\min}(X) \cdot \varepsilon}))$ time.

Our projection operators into the harmonic and boundary space, along with the projection operator of Cohen et al. [10] into the coboundary space, give all the projection operators needed to compute the Hodge decomposition of 1-chains in K .

Our harmonic projection operator is built using an orthonormal approximate harmonic basis (part (i) of Lemma 3). Dey [14] describes a nearly-linear time algorithm for computing a homology basis for a complex linearly embedded in \mathbb{R}^3 . Black et al. [3] describe a nearly-linear time algorithm for computing a cohomology basis for subcomplexes of collapsible complexes embedded in \mathbb{R}^3 . Our harmonic basis, though approximate, can be viewed as a complement to these two results.

1.2 Paper organization

In addition to this introduction, the main body of this paper is a background and overview section. To simplify the presentation, the bulk of the technical details are left for the full version of this paper, and the overview provides a high-level description of our approaches as well as the technical challenges and contribution of this paper. In the overview, we included references to the technical lemmas to enable easy access to the full paper.

The background section introduces standard definitions of the concepts used in this paper. We hope this section provides easy lookup for the reader while reading the overview section as well as the technical part of the paper.

2 Background

In this section, we review basic definitions from linear algebra and algebraic and combinatorial topology that are used in this paper; see references [7, 17, 18, 31] for further background.

2.1 Linear Algebra

Span, Basis. Let $V = \{v_1, \dots, v_k\}$ be a set of vectors in \mathbb{R}^n . The **span** of V , denoted $\text{span}(V)$, is the subspace of \mathbb{R}^n of all linear combinations of V . In particular, V spans \mathbb{R}^n if any vector in \mathbb{R}^n is a linear combination of the vectors in V . We say that V is a **basis** for its span if the dimension of its span equals the cardinality of V .

Linear map, projection, inverse. Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map, represented by an $m \times n$ matrix. Typically, we don't make a distinction between a linear map and its matrix representation and denote both as A . The **kernel** of A is $\ker(A) := \{x \in \mathbb{R}^n : Ax = 0\}$, and the **image** of A is $\text{im}(A) = \{Ax : x \in \mathbb{R}^n\}$. The **rank** of a linear map is the dimension of its image.

We say that U and V **orthogonally decompose** W , denoted $W = U \oplus V$, if (i) any vector in U is orthogonal to any vector in V , and (ii) any vector in $x \in W$ is a unique sum of vectors in $x_U \in U$ and $x_V \in V$, i.e. $x = x_U + x_V$. The **fundamental theorem of linear**

algebra states that $\mathbb{R}^n = \text{im}(A^T) \oplus \text{ker}(A)$ and $\mathbb{R}^m = \text{im}(A) \oplus \text{ker}(A^T)$, where A^T is the transpose of A obtained by flipping A over its diagonal. In particular, if $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is symmetric (i.e., $A = A^T$), then $\mathbb{R}^n = \text{im}(A) \oplus \text{ker}(A)$.

A linear map $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a **projection** if it is the identity for the vectors in its image, or equivalently, $AA = A$. The map A is an **orthogonal projection** if it maps each point of \mathbb{R}^n to its closest point in $\text{im}(A)$, or equivalently, $A^T = A = AA$. Note for any subspace U of \mathbb{R}^n there is a unique orthogonal projection into U , denoted Π_U . If $\{u_1, \dots, u_k\}$ is an orthonormal basis for U , the orthogonal projection into U is the linear map $\Pi_U = \sum_{i=1}^k u_i u_i^T$.

If a linear map $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is bijective, it has a well-defined **inverse** denoted $A^{-1} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ where $Ax = b \iff A^{-1}b = x$. More generally, the **pseudoinverse** of A is the unique linear map $A^+ : \mathbb{R}^m \rightarrow \mathbb{R}^n$ with the following properties: (i) $AA^+A = A$, (ii) $A^+AA^+ = A^+$, (iii) $(AA^+)^T = AA^+$, and (iv) $(A^+A)^T = A^+A$. Admittedly, the definition of the pseudoinverse is not very intuitive. A more intuitive description is that A^+ is the unique linear map with the following properties: (1) A^+ maps any vector $y \in \text{im}(A)$ to the unique vector $x \in \text{im}(A^T)$ such that $Ax = y$, and (2) A^+ maps any vector $y \in \text{ker}(A^T)$ to 0. While it is not true in general that $(A + B)^+ = B^+ + A^+$ for linear maps A and B , this is true if $A^T B = B^T A = 0$; see Campbell [7], Theorem 3.1.1.

Matrix norm, singular values, Loewner order. A symmetric matrix A is **positive semi-definite** if $x^T Ax \geq 0$ for each $x \in \mathbb{R}^n$. The **Loewner Order** is a partial order on the set of $n \times n$ symmetric matrices. For symmetric matrices A and B , we say $A \preceq B$ if $B - A$ is positive semidefinite.

Let $x \in \mathbb{R}^n$. Let p be a positive integer. The **p-norm** of x is $\|x\|_p = (\sum_{i=1}^n |x[i]|^p)^{\frac{1}{p}}$. We use the 1-norm and 2-norm in this paper. An important fact we will use throughout this paper is that $\|x\|_2 \leq \|x\|_1 \leq \sqrt{n}\|x\|_2$. For any norm $\|\cdot\|$ on \mathbb{R}^n , there is an accompanying **operator norm** of a matrix A defined $\|A\| = \max_{x:\|x\|=1} \|Ax\|$, or equivalently, $\|A\| = \max_{x \neq 0} (\|Ax\|/\|x\|)$.

Unless otherwise specified, all norms in this paper will be the 2-norm.

The **singular value decomposition** of $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for $m \geq n$ (resp. $m \leq n$) is a set of n (resp. m) orthonormal vectors $\{u_1, \dots, u_n\} \subset \mathbb{R}^n$ called **left singular vectors**, n (resp. m) orthonormal vectors $\{v_1, \dots, v_n\} \subset \mathbb{R}^n$ called **right singular vectors**, and n (resp. m) real numbers $\{\sigma_1, \dots, \sigma_n\} \subset \mathbb{R}$ called **singular values** such that $A = \sum_{i=1}^n \sigma_i u_i v_i^T$. The **condition number** of a linear map $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is $\kappa(A) = |\sigma_{\max}(A)|/|\sigma_{\min}(A)|$, where $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ are the largest and smallest non-zero singular values of A .

The **eigenvectors** and **eigenvalues** of a matrix $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ are n vectors $\{v_1, \dots, v_n\} \subset \mathbb{R}^n$ and n real numbers $\{\lambda_1, \dots, \lambda_n\}$ such that $Av_i = \lambda_i v_i$. The singular values and right singular vectors (resp. left singular values) of a matrix $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ are the square roots of the eigenvalues and eigenvectors of $A^T A$ (resp. AA^T). If a matrix A is symmetric, the eigenvectors of A are orthogonal, and the eigenvectors and eigenvalues of A are the left and right singular vectors and the singular values.

Determinant, Cramer’s rule, unimodularity. For any $1 \leq i \leq n$, the **determinant** of an $n \times n$ matrix $A = [a_{i,j}]_{1 \leq i,j \leq n}$ can be defined via its **Laplace expansion** as $\det(A) = \sum_{j=1}^n ((-1)^{i+j} \cdot a_{i,j} \cdot \det(A_{i,j}))$, where $A_{i,j}$ is the $(n - 1) \times (n - 1)$ matrix obtained by removing the i th row and j th column of A . It is well known that $\det(A) \neq 0$ if and only if A is bijective. In that case, **Cramer’s rule** give an explicit formula for the solution of the linear system $Ax = b$, which is $x[i] = \det(A_i)/\det(A)$ where A_i is the matrix obtained by replacing the i^{th} column of A with b .

An $n \times n$ matrix A is **unimodular** if $\det(A) \in \{-1, +1\}$. By Cramer's rule, $Ax = b$ has an integer solution if A is unimodular and A and b have integer coefficients. An $n \times m$ matrix B is **totally unimodular** if for any square submatrix A of B , $\det(A) \in \{-1, 0, +1\}$. The 1-boundary matrix of a simplicial complex (defined below) is totally unimodular [26].

2.2 Topology

Simplicial complexes. A **simplicial complex** K is a set of finite sets such that if $\tau \in K$ and $\sigma \subset \tau$, then $\sigma \in K$. A **subcomplex** of K is a subset $L \subset K$ such that L is a simplicial complex. The **vertices** of K is the set $\cup_{\sigma \in K} \sigma$. We assume there is a fixed but arbitrary order (v_1, \dots, v_n) on the vertices of K .

An element $\sigma \in K$ with $|\sigma| = d + 1$ is a **d-simplex**. A 0-simplex is a **vertex**, a 1-simplex is an **edge**, a 2-simplex is a **triangle**, and a 3-simplex is a **tetrahedron**. The set of all d -simplices in K is denoted K_d . For two simplices $\tau \subset \sigma$, we say that τ is a **face** of σ .

Hodge decomposition, homology, cohomology. The **dth chain group** of a simplicial complex K is the vector space $C_d(K)$ over \mathbb{R} with orthonormal basis K_d , and an element of $C_d(K)$ is a **d-chain**. The **dth boundary map** is the linear map $\partial_d : C_d(K) \rightarrow C_{d-1}(K)$ defined $\partial_d \sigma = \sum_{i=0}^d (-1)^i (\sigma \setminus \{v_{k_i}\})$ for each simplex $\sigma = \{v_{k_0}, \dots, v_{k_d}\} \in K_d$, where we assume $v_{k_i} < v_{k_j}$ for $i < j$. The **dth coboundary map** is $\partial_{d+1}^T : C_d(K) \rightarrow C_{d+1}(K)$. Elements of $\ker \partial_d$ (resp. $\ker \partial_{d+1}^T$) are **cycles** (resp. **cocycles**), and elements of $\text{im } \partial_{d+1}$ (resp. $\text{im } \partial_d^T$) are **boundaries** or **null-homologous cycles** (resp. **coboundaries**.) Two cycles (resp. cocycles) γ_1 and γ_2 are **homologous** (resp. **cohomologous**) if their difference $\gamma_1 - \gamma_2$ is a boundary (resp. coboundary.)

The **dth Laplacian** is the linear map $L_d : C_d(K) \rightarrow C_d(K)$ defined $L_d = \partial_d^T \partial_d + \partial_{d+1} \partial_{d+1}^T$. The **dth up-Laplacian** is the linear map $L_d^{up} = \partial_{d+1} \partial_{d+1}^T$, and the **dth down-Laplacian** is the linear map $L_d^{down} = \partial_d^T \partial_d$.

A key fact of algebraic topology is that $\partial_d \partial_{d+1} = 0$, hence $\text{im } \partial_{d+1} \subset \ker \partial_d$, and $\text{im } \partial_d^T \subset \ker \partial_{d+1}^T$. The **dth homology group** is the quotient group $H_d(K) = \ker \partial_d / \text{im } \partial_{d+1}$, and the **dth cohomology group** is the quotient group $H^d(K) = \ker \partial_{d+1}^T / \text{im } \partial_d^T$. Since $\text{im } \partial_d^T \oplus \ker \partial_d$ and $\text{im } \partial_{d+1} \oplus \ker \partial_{d+1}^T$ are two orthogonal decompositions of the d -chain space, the d^{th} homology group and the d^{th} cohomology group have the same rank, which is the **dth Betti number** of the complex, denoted $\beta_d(K)$. We say two cycles are **homologous** (resp. **cohomologous**) if they are in the same homology (resp. cohomology) class, or equivalently, if their difference is a boundary (resp. coboundary.)

The **Hodge Decomposition** is the orthogonal decomposition of the d^{th} chain group into $C_d(K) = \text{im}(\partial_{d+1}) \oplus \ker(L_d) \oplus \text{im}(\partial_d^T)$. The subspace $\ker(L_d)$ are the **harmonic chains**. Thus, any chain $x \in C_d(K)$ can be uniquely written as the sum $x = x_{bd} + x_{hr} + x_{cbd}$ where $x_{bd} \in \text{im}(\partial_{d+1})$, $x_{hr} \in \ker(L_d)$, and $x_{cbd} \in \text{im}(\partial_d^T)$.

A **d-boundary basis**, **d-coboundary basis** and **d-harmonic basis** are bases for the boundary, coboundary and harmonic spaces. A **d-homology basis** is a maximal set of cycles such that no linear combination of these cycles is a boundary. Similarly, a **d-cohomology basis** is a maximal set of cocycles such that no linear combination of these cocycles is a coboundary. We have the following fact.

► **Fact 4.** *A set of cycles (resp. cocycles) is a homology (resp. cohomology) basis if and only if their projection into the harmonic space is a harmonic basis.*

Two cycles (resp. cocycles) are homologous (resp. cohomologous) if they have the same harmonic part, as then their difference is a boundary (resp. coboundary). Accordingly, the previous fact implies that for any cycle (resp. cocycle) x and any homology basis

(resp. cohomology basis) Γ , there is a unique linear combination of the elements of Γ that is homologous (resp. cohomologous) to x ; this is the linear combination of Γ with the same harmonic component as x .

A useful property of cohomology bases is they can be used to tell when two cycles are homologous, as described by the following fact.

► **Fact 5** (Busaryev et al. [6]). *Let x and y be cycles (resp. cocycles), and let P be a cohomology basis (resp. homology basis.) Then y is homologous (resp. cohomologous) to x if and only if $x \cdot p = y \cdot p$ for all $p \in P$.*

Collapsibility. Let K be a simplicial complex, σ a d -simplex of K , and τ a $(d-1)$ -simplex of K that is a face of σ . If τ is not the face of any other simplex, we say that K **collapses** into $K \setminus \{\sigma, \tau\}$; we refer to (σ, τ) as a **collapse pair**. Moreover, we say that a complex collapses to itself. Inductively, we say that a complex K **collapses** into a complex K' if there is a complex K'' such that K collapses to K'' and K'' collapses to K' . We say that a complex K is **collapsible** if it collapses to a single vertex.

When a complex K collapses to a complex K' , we obtain a sequence of complexes $K = K_0 \supset K_1 \supset \dots \supset K_t = K'$, where for each $1 \leq i \leq t$, K_i can be obtained from K_{i-1} by removing one collapse pair. We refer to this sequence as a **collapsing sequence**. The complexes K and K' are homotopy equivalent if one collapses to the other, thus, K and K' have isomorphic homology group. In particular, a collapsible complex has trivial homology groups in every nonzero dimension.

Embeddability. A d -dimensional simplicial complex K is **embedded** if $K \subset R$ for R a triangulation of \mathbb{R}^{d+1} . Furthermore, K is **linearly embedded** if there is a homeomorphism from the underlying space $|R|$ to \mathbb{R}^{d+1} that is linear on each simplex, i.e. each 1-simplex is mapped to a line segment, each 2-simplex is mapped to a triangle, etc. All embedded complexes in this paper will be linearly embedded.

We will make use of the dual graph of an embedded complex. Informally, the **dual graph** of an embedded complex is the graph K^* with vertices that are the connected components of $R \setminus K$ and edges between two vertices if there is a d -simplex in K incident to both connected components. Alternatively, the dual graph can be defined with vertices corresponding to a generating set of d -cycles of K . For this construction, see the definition of **Lefschetz set** in the paper [3].


3 Overview

Let X be a collapsible simplicial complex embedded in \mathbb{R}^3 , and let $K \subseteq X$ be a subcomplex of X . We study two closely related problems: (i) computing the Hodge decomposition of the 1-chains of K , and (ii) solving a linear system $L_1 x = b$, where L_1 is the 1-Laplacian of K (in the overview section, all the operators are with respect to K unless mentioned otherwise.) These two problems are related, as our approximate Laplacian solver uses an approximate Hodge decomposition of the input vector x . More generally, understanding the Hodge decomposition is key to understanding this paper as many proofs rely on some property of the Hodge decomposition. Therefore, we begin our overview with an introduction to the Hodge decomposition.

3.1 The Hodge Decomposition

The Hodge decomposition is a decomposition of the chain group $C_d(K)$ in terms of the kernels and images of the boundary operators ∂_d and ∂_{d+1} and their transposes. Specifically, the problems in this paper consider the first chain group $C_1(K)$, the two boundary operators ∂_2 and ∂_1 , and their corresponding coboundary operators ∂_2^T and ∂_1^T . The boundary operator ∂_2 maps each (oriented) triangle to the edges in its boundary; similarly, ∂_1 maps each edge to its two endpoints. A key fact is that $\partial_1\partial_2 = 0$, or equivalently, $\text{im}(\partial_2) \subseteq \text{ker}(\partial_1)$. This implies $\text{im}(\partial_2)$ is orthogonal to $\text{im}(\partial_1^T)$. The former subspace $\text{im}(\partial_2)$ is called the **boundary subspace**, and the latter subspace $\text{im}(\partial_1^T)$ is called the **coboundary subspace**. If K has trivial 1-homology, then $\text{im}(\partial_2) = \text{ker}(\partial_1)$, and the boundary and coboundary spaces give a full orthogonal decomposition of $C_1(K)$, called the **Helmholtz decomposition**. Otherwise, there is a third subspace orthogonal to both the boundary and coboundary subspaces, called the **harmonic subspace**. The harmonic subspace is exactly $\text{ker}(L_1) = \text{ker}(\partial_1) \cap \text{ker}(\partial_2^T)$.

The boundary, coboundary, and harmonic subspaces give a full orthogonal decomposition of $C_1(K)$ called the **Hodge decomposition**, which generalizes the Helmholtz decomposition. Thus, we can express any 1-chain x as $x = x_{cbd} + x_{bd} + x_{hr}$, where x_{cbd} , x_{bd} and x_{hr} are the coboundary, boundary and harmonic part of x and are pairwise orthogonal. The chains $x_{bd} + x_{hr}$ and $x_{cbd} + x_{hr}$ are called the **cyclic** and **cocyclic** parts of x respectively. Similarly, the space spanned by harmonic and boundary chains is called the **cycle space**, and the space spanned by harmonic and coboundary chains is called the **cocycle space**. It is implied by $\partial_1\partial_2 = 0$ that the cycle space and cocycle space are the kernels of ∂_1 and ∂_2^T , respectively. The following figure is an illustration of the Hodge decomposition. Boundary, coboundary, harmonic, cycle, and cocycle spaces are shown using the abbreviations bd, cbd, hr, cyc, and cocyc respectively.



To compute the Hodge decomposition, one seeks orthogonal projection operators into the coboundary, boundary and harmonic subspaces. Let Π_{cbd} , Π_{bd} , and Π_{hr} denote these projection operators. Cohen et al. show that for any 1-chain x , its projection into the coboundary space, $\Pi_{cbd}x$, and cycle space, $\Pi_{cyc}x$, can be approximated quickly with operators $\tilde{\Pi}_{cbd}$ and $\tilde{\Pi}_{cyc}$. These projection operators are a key ingredient of their 1-Laplacian solver, as well as the more recent 1-Laplacian solver described by Black et al.; however, both papers are restricted to cases where the first homology group $H_1(K) = 0$. In this paper, we show that for any x , its projection into the boundary space, $\Pi_{bd}x$, can also be approximated quickly. This new projection operator will allow us to generalize the 1-Laplacian solver of Black et al. to complexes with arbitrary first homology. We also give an approximate projection operator into the harmonic space, but our approximation guarantee for this projection operator is weaker (more below).

3.2 Laplacian Solvers

The 1-Laplacian matrix is defined $L_1 = \partial_2 \partial_2^T + \partial_1^T \partial_1$. To solve a linear system $L_1 x = b$, one seeks to approximate L_1^+ , the pseudoinverse of L_1 . As the images of $\partial_2 \partial_2^T$ and $\partial_1^T \partial_1$ are orthogonal, then $L_1^+ = (\partial_2 \partial_2^T)^+ + (\partial_1^T \partial_1)^+$ (see Campbell [7, Theorem 3.1.1]). Therefore, one can approximate L_1^+ by approximating $(\partial_2 \partial_2^T)^+$ and $(\partial_1^T \partial_1)^+$ individually. Computing $(\partial_1^T \partial_1)^+$ is purely a graph problem as ∂_1 is only defined with respect to the vertices and edges of a complex. Cohen et al. show how to approximate $(\partial_1^T \partial_1)^+$ for general complexes [10, Lemma 3.2]. Approximating $(\partial_2 \partial_2^T)^+$ is a more challenging problem that requires taking into account the relationship between triangles and the edges. Our algorithm for approximating $(\partial_2 \partial_2^T)^+$ relies on our new boundary projection operator, the collapsibility of X , and the embedding of X in \mathbb{R}^3 .

Cohen et al. show how to approximate $(\partial_2 \partial_2^T)^+$ for collapsible complexes embedded in \mathbb{R}^3 . Black et al. generalize their work to obtain an approximate solver for a subcomplex of a collapsible complex in \mathbb{R}^3 provided the subcomplex has trivial homology. Their solver is based on the following general lemma regarding approximations of $(BB^T)^+$ for a general matrix B .

► **Lemma 6** (Black et al. [3], Lemma 4.1). *Let B be a linear operator, let $0 \leq \varepsilon < 1$, and let $\tilde{\Pi}_{\text{im}(B)}$ and $\tilde{\Pi}_{\text{ker}^\perp(B)}$ be symmetric matrices such that $(1 - \varepsilon)\Pi_{\text{im}(B)} \preceq \tilde{\Pi}_{\text{im}(B)} \preceq \Pi_{\text{im}(B)}$, and $(1 - \varepsilon)\Pi_{\text{ker}^\perp(B)} \preceq \tilde{\Pi}_{\text{ker}^\perp(B)} \preceq \Pi_{\text{ker}^\perp(B)}$. Also, let U be a linear map such that for any $y \in \text{im}(B)$, $BUy = y$. We have*

$$(1 - (2\kappa + 1)\varepsilon)(BB^T)^+ \preceq \tilde{\Pi}_{\text{im}(B)} U^T \tilde{\Pi}_{\text{ker}^\perp(B)} U \tilde{\Pi}_{\text{im}(B)} \preceq (1 + \kappa\varepsilon)(BB^T)^+,$$

where κ is the condition number of BB^T within the image of B .

This lemma shows the following linear operators are sufficient for approximating $(\partial_2 \partial_2^T)^+$.

- (i) An operator U that for 1-boundary $y \in \text{im}(\partial_2)$ returns a 2-chain $x = Uy$ such that $\partial_2 x = y$. For other vectors $z \notin \text{im}(\partial_2)$, U can return anything as long as U is still linear.
- (ii) An approximate orthogonal projection operator into $\text{im}(\partial_2^T)$, the coboundary space of 2-chains.
- (iii) An approximate orthogonal projection operator into $\text{im}(\partial_2)$, the boundary space of 1-chains.

Black et al. describe an algorithm for computing U that uses the collapsibility and embedding of the supercomplex X . Cohen et al. show that the 2-coboundary space of embedded complexes is dual to the 1-cycle space of the dual graph, hence projection into this space can be approximated using $\tilde{\Pi}_{\text{cyc}}$. Finally, lacking an approximate projection into the boundary space of 1-chains, they needed to assume that their complex has trivial first homology (i.e. that $\text{im}(\partial_2) = \text{ker}(\partial_1)$) so that they can instead use the projection operator into the cycle space of Cohen et al. The boundary projection operator described in this paper allow us to remove that assumption to obtain a solver for any subcomplex K of X . The running time of our new solver polynomially depends on the rank of the homology group and nearly-linearly depends on the size of the complex. We give a complete analysis of our solver in Appendix D of the full paper.

In the rest of this section, we sketch the high level ideas for computing our approximate projection operators. But before we can do that, we need to explain the two notions of approximations that are used in this paper.

3.3 Loewner order approximation

We use the Loewner order on positive semidefinite matrices to specify the approximation quality of our projection and pseudoinverse operators. We see two types of approximation guarantees in this paper for an operator A : **input-relative error bounds** of the form $-\varepsilon I \preceq A - \tilde{A} \preceq \varepsilon I$ and **output-relative error bounds** of the form $-\varepsilon A \preceq A - \tilde{A} \preceq \varepsilon A$. Note for any vector x , an input relative error bound implies $\|(A - \tilde{A})x\| \leq \varepsilon\|x\|$ – the error is bounded relative to the size of the input x – while an output-relative error bound implies $\|(A - \tilde{A})x\| \leq \varepsilon\|Ax\|$ – the error is bounded relative to the size of the output Ax . An approximate operator with a small input-relative error can have arbitrarily large output-relative error, for example when x is in the kernel of A . Further, output-relative error bounds are stronger if the norm of $\|A\|$ is at most one, i.e. $\|Ax\| \leq \|x\|$, which is the case for the orthogonal projection operators of the Hodge decomposition.

We achieve an output-relative error bound for our approximation of $(L_1[K])^+$. Further, we achieve an output-relative error bound for our approximation $\tilde{\Pi}_{bd}$ of Π_{bd} , but an input-error bound for our approximation $\tilde{\Pi}_{hr}$ of Π_{hr} :

$$-\varepsilon I \preceq -\varepsilon \Pi_{bd} \preceq \Pi_{bd} - \tilde{\Pi}_{bd}(\varepsilon) \preceq \varepsilon \Pi_{bd} \preceq \varepsilon I, \quad (1)$$

and

$$-\varepsilon I \preceq \Pi_{hr} - \tilde{\Pi}_{hr}(\varepsilon) \preceq \varepsilon I. \quad (2)$$

Previously, Cohen et al. had shown approximation operators $\tilde{\Pi}_{cbd}$ and $\tilde{\Pi}_{cyc}$ for projecting into the coboundary and cycle spaces with output-relative error bounds:

$$-\varepsilon I \preceq -\varepsilon \Pi_{cbd} \preceq \Pi_{cbd} - \tilde{\Pi}_{cbd}(\varepsilon) \preceq \varepsilon \Pi_{cbd} \preceq \varepsilon I, \quad (3)$$

and

$$-\varepsilon I \preceq -\varepsilon \Pi_{cyc} \preceq \Pi_{cyc} - \tilde{\Pi}_{cyc}(\varepsilon) \preceq \varepsilon \Pi_{cyc} \preceq \varepsilon I. \quad (4)$$

We use these operators multiple times in our algorithms. For simplification, we drop the explicit mention of the parameter ε when it is clear from the context.

3.4 Projection operators

We first describe our algorithm for computing $\tilde{\Pi}_{hr}$ (an overview of Appendices A and B of the full version of the paper). Based on that and the operator $\tilde{\Pi}_{cbd}$ of Equation (3), we show how to compute $\tilde{\Pi}_{bd}$ (an overview of Appendix C of the full version of the paper).


3.4.1 Harmonic projection

We compute our approximate harmonic projection operator $\tilde{\Pi}_{hr}$ by computing an approximate orthonormal basis $\tilde{G} = \{\tilde{g}_1, \dots, \tilde{g}_\beta\}$ of the harmonic space. We then define the approximate projection into the harmonic space to be the linear map $\tilde{\Pi}_{hr} = \sum_{i=1}^{\beta} \tilde{g}_i \tilde{g}_i^T$.

To compute \tilde{G} , our algorithm starts with a cohomology basis $P = \{p_1, \dots, p_\beta\}$; the algorithm for computing P is given at the end of this section. From P , it computes $\tilde{H} = \{\tilde{h}_1, \dots, \tilde{h}_\beta\}$, where $\tilde{h}_i = \tilde{\Pi}_{cyc} p_i$ and $\tilde{\Pi}_{cyc}$ is the projection operator of Equation (4). The set \tilde{H} is an approximate harmonic basis, but it is not orthonormal. Next, we normalize \tilde{H} to obtain $\tilde{N} = \{\tilde{h}_1/\|\tilde{h}_1\|, \dots, \tilde{h}_\beta/\|\tilde{h}_\beta\|\}$. Finally, we run Gram-Schmidt on \tilde{N} to obtain \tilde{G} .

To see why \tilde{G} is an approximate basis for the harmonic space, let us consider a much easier analysis assuming we can use the exact projection in the cycle space Π_{cyc} instead of the approximate projection $\tilde{\Pi}_{cyc}$. Instead of \tilde{H} , \tilde{N} and \tilde{G} , let $H = \{h_1, \dots, h_\beta\}$, $N = \{h_1/\|h_1\|, \dots, h_\beta/\|h_\beta\|\}$ and $G = \{g_1, \dots, g_\beta\}$ be the sets of vectors we obtain when we use the exact projection operator. In that case, $h_i = \Pi_{cyc} p_i$ is the harmonic part of p_i ; this is because p_i is a cocycle, so projecting it into the cycle space is the same as projecting it into the harmonic space. It follows from Fact 4 in Section 2 that G is an exact orthonormal basis of the harmonic space, thus it defines an exact projection operator into the harmonic space.

In the real scenario where we work with the approximate projection operator $\tilde{\Pi}_{cyc}$, two undesirable things can happen. First, we can no longer guarantee that the vectors in \tilde{N} are purely harmonic, as the error introduced by the approximate operator $\tilde{\Pi}_{cyc}$ may be part boundary. However, this is not an issue, as we can make the boundary components of \tilde{N} sufficiently small by approximating $\tilde{\Pi}_{cyc}$ more accurately. Second, and more importantly for our application, the spaces spanned by N and \tilde{N} can be very different, even if the vectors N and \tilde{N} are pairwise close. As an example, imagine that we have two pairs of vectors $N = \{\eta_1, \eta_2\}$ and $\tilde{N} = \{\tilde{\eta}_1, \tilde{\eta}_2\}$ such that $\|\eta_i - \tilde{\eta}_i\| < \varepsilon$ for $i = 1, 2$. We might guess that the two spaces spanned by N and \tilde{N} are similar as the vector are close, but if η_1 and η_2 are also close, then the two vectors spaces can be drastically different. Figure 1 gives an illustration of this, where N is the set of blue vectors and \tilde{N} is the set of red vectors. As illustrated in the figure, the space spanned by N and the space spanned by \tilde{N} can be drastically different.



■ **Figure 1** Pairwise closeness between a set of vectors N and \tilde{N} is not enough to guarantee the spaces spanned by N and \tilde{N} are close! The red and blue vectors are pairwise close, but the spaces they span are very different.

We can remedy this if we approximate \tilde{N} within a sufficiently small error ε of N , but this new error bound needs to take into account the similarity of the vectors in N . The question is how accurately we need to approximate $\tilde{\Pi}_{cyc}$ to obtain a sufficiently small approximation error for \tilde{N} . To answer this question, we define a measure of linear independence of N called its δ -independence. Formally, we say that N is **δ -independent** if each vector $h_i/\|h_i\| \in N$ is at distance at least δ from the span of the other vectors of N . Intuitively, larger δ means N is more independent, in the sense that the elements are well-separated. The smaller the δ , the more accurately we need to approximate $\tilde{\Pi}_{cyc}$ to ensure that N and \tilde{N} will span similar spaces. This intuition is summarized by the following lemma, showing the error in projection into N as a function of δ , ε , and β , where ε bounds the difference between N and \tilde{N} .

► **Corollary 7.** *Let $0 < \delta < 1$, and let $0 < \varepsilon < \left(\frac{\delta}{8\beta}\right)^\beta$. Let $N = \{\eta_1, \dots, \eta_\beta\}$ be a set of δ -linearly independent unit vectors, and let $\tilde{N} = \{\tilde{\eta}_1, \dots, \tilde{\eta}_\beta\}$ be a set of unit vectors such that $\|\eta_i - \tilde{\eta}_i\| < \varepsilon$. Let $G = \{g_1, \dots, g_\beta\}$ be the orthonormal basis that is the output of running Gram-Schmidt on N , and let $\tilde{G} = \{\tilde{g}_1, \dots, \tilde{g}_\beta\}$ be the output of running Gram-Schmidt on \tilde{N} . Then $\left\| \Pi_{\text{span } N} - \Pi_{\text{span } \tilde{N}} \right\| = \left\| \sum_{i=1}^\beta g_i g_i^T - \tilde{g}_i \tilde{g}_i^T \right\| < 2 \cdot \beta \cdot \left(\frac{8\beta}{\delta}\right)^\beta \varepsilon$.*

The difficulty here is actually determining a lower bound on the δ -independence of N . We have access to the cohomology basis P , but we need the (normalized) harmonic parts of P to be δ -independent. Note that P can be composed of vectors that are very strongly independent, yet their harmonic parts may only be weakly independent, for example, when the vectors of P have similar harmonic parts but very different coboundary parts.

We show that if P is composed of integer vectors with maximum length p_{\max} , then P being linearly independent implies that H is δ -independent for a $\delta \sim 1/(p_{\max} \cdot n_1)^\beta$, where n_1 is the number of edges in K . In addition to the properties of P , our proof of Lemma 8 relies on the total unimodularity of ∂_1 .

► **Lemma 8.** *Let K be a simplicial complex with n_1 edges such that $H_1(K) = \beta$. Let $\{p_1, \dots, p_\beta\}$ be a 1-cohomology basis for K such that each p_i is an integer vector with maximum Euclidean norm p_{\max} . Let h_i be the harmonic part of p_i for $1 \leq i \leq \beta$. Then*

- (i) $\|h_i\| \geq 1/(\sqrt{n_1} \cdot p_{\max})^\beta$ for each $1 \leq i \leq \beta$, and
- (ii) $\{h_1/\|h_1\|, \dots, h_\beta/\|h_\beta\|\}$ is $(1/(\sqrt{n_1} \cdot p_{\max})^\beta)$ -independent.

Computing a Cohomology Basis. The question remains of how to find the cohomology basis P . For this, we use an algorithm on Black et al. Dey [14] describes a nearly-linear time algorithm for computing a homology basis composed of vectors with coordinates in $\{-1, 0, +1\}$. Black et al. [3] describe an operator $C(X, K)$ that when applied to a homology basis returns a cohomology basis. We use the cohomology basis P obtained by applying the operator C to Dey's homology basis.

► **Lemma 9** (Dey, Lemma [14]). *For a 2-dimensional simplicial complex K linearly embedded in \mathbb{R}^3 , there exists an algorithm computing a basis for $H_1(K, \mathbb{R})$ in $O(n \log n + n\beta_1)$ time, where n is the complexity of K . Further, the basis is composed of vectors with all coordinates from $\{-1, 0, +1\}$.*

► **Lemma 10** (Black et al., Lemma 1.1 [3]). *Let X be a collapsible simplicial complex in \mathbb{R}^3 , and let $K \subset X$ be a subcomplex of X . Let β be the rank of $H_1(K)$ and let n be the total number of simplices of X . Let $\Gamma = \{\gamma_1, \dots, \gamma_\beta\}$ be a homology basis for K . There is a linear operator $C(X, K)$ such that the set $C_\Gamma = \{C(X, K) \cdot \gamma_1, \dots, C(X, K) \cdot \gamma_\beta\}$ is a cohomology basis for K . Furthermore, C_Γ can be computed in $O(\beta \cdot n)$ time.*

While Black et al. introduced the algorithm for computing the cohomology basis P , we bound the length of the vectors in P . We bound the length of P by combining a bound on the length of the homology basis with the following bound on the operator norm $\|C(X, K)\|$.

► **Lemma 11.** *There is a constant α such that $\|C(X, K)\| \leq \alpha \cdot n_1^2 n_2^4 / (\lambda_{\min}(L_1^{up}[X]))$, where $\lambda_{\min}(L_1(X))$ is the smallest non-zero eigenvalue of $L_1(X)$ and n_1 and n_2 are the number of edges and triangles in X respectively.*

Note that the vectors in Dey's homology basis have bounded length as they have coefficients $\{-1, 0, 1\}$. By combining this observation, Lemma 8, and Lemma 11, we prove that the harmonic part of P are δ -independent for an appropriate value of δ . This is summarized in the following corollary.

► **Corollary 12.** *Let X be a collapsible complex embedded in \mathbb{R}^3 with a known collapsing sequence and let $K \subset X$ be a subcomplex of X . Let β be the rank of $H_1(K)$ and let n be the total number of simplices in X . There is an $O(n \log n + \beta n)$ time algorithm for computing a cohomology basis $\{p_1, \dots, p_\beta\}$ of K with harmonic parts $\{h_1, \dots, h_\beta\}$ such that*

- (i) *each h_i has length at least δ , and*
 - (ii) *the set $\{h_1/\|h_1\|, \dots, h_\beta/\|h_\beta\|\}$ is δ -linearly independent,*
- where $\delta = (\lambda_{\min}(L_2^{up}(X)))/(\alpha \cdot n_1^3 n_2^4)^\beta$ for a constant α .

Passing this cohomology basis P to the algorithm above, we obtain \tilde{H} , \tilde{N} , \tilde{G} , $\tilde{\Pi}_{hr}$ as desired. The exact approximation quality of the approximate harmonic basis and approximate harmonic projection are given in Lemma 3 in the introduction.

3.4.2 Boundary projection

It follows from the Hodge decomposition that the the projection into the boundary space can be written $\Pi_{bd} = I - \Pi_{cbd} - \Pi_{hr}$. We have approximate projections $\tilde{\Pi}_{cbd}$ and $\tilde{\Pi}_{hr}$ with input-relative error bounds (Equations (3) and (2) respectively), so we immediately obtain a boundary projection $\bar{\Pi}_{bd}$ with input-relative error bound defined

$$\bar{\Pi}_{bd} = I - \tilde{\Pi}_{cbd} - \tilde{\Pi}_{hr} \implies -\varepsilon I \preceq \Pi_{bd} - \bar{\Pi}_{bd} \preceq \varepsilon I.$$

However, we need a boundary projection operator with an output-relative bound for our solver. Unfortunately, the operator $\bar{\Pi}_{bd}$ can have arbitrarily bad output-relative error. Specifically, for any vector x that is orthogonal to the boundary space, this operator has unbounded output-relative error as $\bar{\Pi}_{bd}x = 0$.

We instead use $\bar{\Pi}_{bd}$ as a starting point for a projection operator with bounded output-relative error. To that end, let's revisit the issue of input vs. output relative error. Let $x = x_{bd} + x_{cocyc}$ be any vector decomposed into its boundary and cocycle parts. The input-relative error bound of $\bar{\Pi}_{bd}$ is proportional to $\|x\| = \|x_{bd} + x_{cocyc}\|$, while for an output-relative error bound, we need the bound to be proportional to $\|\Pi_{bd}x\| = \|x_{bd}\|$. Therefore, a problem arises if x_{cocyc} is much larger than x_{bd} ; provided a bound on $\|x_{cocyc}\|/\|x_{bd}\|$, we can accordingly modify the accuracy of our projection operators $\tilde{\Pi}_{cbd}(\varepsilon)$ and $\tilde{\Pi}_{hr}(\varepsilon)$ to ensure $\bar{\Pi}_{bd}$ has small output-relative error for x . Unfortunately, $\|x_{cocyc}\|/\|x_{bd}\|$ can be unbounded. To counteract this, we show that we can map x to a different vector x' before passing it to $\bar{\Pi}_{bd}$ such that (1) x' has the same boundary component as x (so $\Pi_{bd} \cdot x = \Pi_{bd} \cdot x'$), and (2) $\|x'_{cocyc}\|/\|x'_{bd}\|$ is bounded.

Specifically, our boundary projection operator is defined $\tilde{\Pi}_{bd} = (I - P_\Gamma)(I - P_T)\bar{\Pi}_{bd}(I - P_T)^T(I - P_\Gamma)^T$, defined based on two operators P_T and P_Γ . The former was introduced by Cohen et al. to obtain $\tilde{\Pi}_{cyc}$, and the latter is introduced in this paper; we sketch the ideas of both in this overview. The operator $(I - P_T)^T(I - P_\Gamma)^T$ behaves as we need: it maps x to a chain x' with the same boundary component as x and a relatively bounded cocycle part. We now describe P_T and P_Γ .

Let T be any spanning tree of the 1-skeleton of K . P_T is the operator that maps any 1-chain to the unique 1-chain with the same boundary in T . In particular, for any 1-chain x , $(I - P_T)x$ is a cycle.

Next, let $\Gamma = \{\gamma_1, \dots, \gamma_\beta\}$ be a 1-homology basis in K . P_Γ is the operator that maps any 1-cycle to the unique linear combination of Γ that is in the same homology class. In particular, for any 1-cycle x , $(I - P_\Gamma)x$ is a boundary.

Now let $F = (I - P_T)^T(I - P_\Gamma)^T$, so $F^T = (I - P_\Gamma)(I - P_T)$. Consider any vector $x = x_{bd} + x_{cbd} + x_{hr}$. We investigate what F does to each of the three constituents of x ; what can we say about Fx_{bd} , Fx_{cbd} and Fx_{hr} ? In what follows, we frequently use the fact that for any linear map A , $\ker(A)$ and $\text{im}(A^T)$ orthogonally decompose the domain of A .

$(I - P_T)$ maps any 1-chain to a cycle and $(I - P_\Gamma)$ maps any cycle to a boundary cycle; thus, $\text{im}(F^T)$ is a subset of the boundary space. It follows that $\ker(F)$ is a superset of the orthogonal complement of the boundary space, which is the cocycle space. So, F maps any cocycle to zero, in particular, $Fx_{cbd} = 0$ and $Fx_{hr} = 0$. It remains to investigate Fx_{bd} .

P_T maps any cycle to zero, so $\ker(P_T)$ includes the cycle space; hence, $\text{im}(P_T^T)$ is a subset of the orthogonal complement of the cycle space, which is the coboundary space. In particular, $\text{im}(P_T^T)$ is a subset of the cocycle space. In addition, P_Γ maps all boundary cycles to zero, so $\ker(P_\Gamma)$ includes the boundary space; hence, $\text{im}(P_\Gamma^T)$ is within the orthogonal complement of the boundary space, which is the cocycle space. Now consider

$$Fx_{bd} = (I - P_T)^T(I - P_\Gamma)^T x_{bd} = Ix_{bd} - P_T^T(I - P_\Gamma)^T x_{bd} - P_\Gamma^T x_{bd} = x_{bd} + x'_{cocyc},$$

and observe that x'_{cocyc} is indeed in the cocycle space as $\text{im}(P_T^T)$ and $\text{im}(P_\Gamma^T)$ are both within this space.

Overall, $Fx_{cbd} = 0$, $Fx_{hr} = 0$ and $Fx_{bd} = x_{bd} + x'_{cocyc}$, so Fx has the same boundary part as x . Moreover, the norm of the cocyclic part of Fx , $\|x'_{cocyc}\|$, can now be bounded by $\|F\| \cdot \|x_{bd}\|$, as it is produced by applying F to x . The proof of Lemma 3.2 in Cohen et al. and Corollary 14 of this paper provide bounds on $\|I - P_T\|$ and $\|I - P_\Gamma\|$ that are dependent on the number of simplices of X , the smallest non-zero eigenvalue of the up-Laplacian of X , and the first Betti number of K ; these bounds in turn provide a bound on $\|F\|$. (Note that $\|A\| = \sqrt{\|AA^T\|}$ for any linear operator A .)

► **Lemma 13** (Cohen et al., Proof of Lemma 3.2 [10]). *Let K be any simplicial complex, and let T be a spanning tree of the 1-skeleton of K . Let P_T be the operator that maps any 1-chain x to the unique 1-chain on T with the same boundary, that is (i) $P_T \cdot x \in C_1(T)$, and (ii) $\partial_1 x = \partial_1 P_T \cdot x$. We have $(I - P_T)(I - P_T)^T \preceq n_1^2 I$, where n_1 is the number of edges of K . Further, for any x , $P_T \cdot x$ can be computed in $O(n_1)$ time.*

► **Corollary 14.** *Let X, K as defined. Let $\Gamma = \{\gamma_1, \dots, \gamma_\beta\}$ be the homology basis of Lemma 9, and let P_Γ be the operator that for any cycle α returns the unique linear combination of the cycles of Γ that is homologous to α . We have*

$$(1 - P_\Gamma)(1 - P_\Gamma)^T \preceq \varepsilon \cdot I,$$

for $\varepsilon = (n_1 n_2 / \lambda_{\min}(X))^{c \cdot \beta}$, where $\lambda_{\min}(X)$ is the smallest non-zero eigenvalue of $L_1^{up}(X)$ and c is a constant. Further, for any vector v , $P_\Gamma \cdot v$ can be computed in $O(\beta^2 n_1 + \beta^\omega)$ time.

The accuracy and time complexity of the approximate boundary solver $\tilde{\Pi}_{bd}$ are described in Lemma 2 in the introduction.

References

- 1 Douglas Arnold, Richard Falk, and Ragnar Winther. Finite element exterior calculus: from hodge theory to numerical stability. *Bulletin of the American Mathematical Society*, 47(2):281–354, January 2010. doi:10.1090/s0273-0979-10-01278-4.
- 2 Marshall Bern, John R. Gilbert, Bruce Hendrickson, Nhat Nguyen, and Sivan Toledo. Support-graph preconditioners. *SIAM J. Matrix Anal. Appl.*, 27(4):930–951, December 2005. doi:10.1137/S0895479801384019.

- 3 Mitchell Black, William Maxwell, Amir Nayyeri, and Eli Winkelman. Computational topology in a collapsing universe: Laplacians, homology, cohomology. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 226–251, 2022. doi:10.1137/1.9781611977073.12.
- 4 Erik G. Boman and Bruce Hendrickson. Support theory for preconditioning. *SIAM J. Matrix Anal. Appl.*, 25(3):694–717, March 2003. doi:10.1137/S0895479801390637.
- 5 Erik G. Boman, Bruce Hendrickson, and Stephen Vavasis. Solving elliptic finite element systems in near-linear time with support preconditioners. *SIAM J. Numer. Anal.*, 46(6):3264–3284, October 2008. doi:10.1137/040611781.
- 6 Oleksiy Busaryev, Sergio Cabello, Chao Chen, Tamal K. Dey, and Yusu Wang. Annotating simplices with a homology basis and its applications. In Fedor V. Fomin and Petteri Kaski, editors, *Algorithm Theory – SWAT 2012*, pages 189–200, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 7 S. L. (Stephen La Vern) Campbell. *Generalized inverses of linear transformations*. Surveys and reference works in mathematics. Pitman, London, 1979.
- 8 Ozan Candogan, Ishai Menache, Asuman E. Ozdaglar, and Pablo A. Parrilo. Flows and decompositions of games: harmonic and potential games. *Math. Oper. Res.*, 36(3):474–503, 2011. doi:10.1287/moor.1110.0500.
- 9 Paul Christiano, Jonathan A. Kelner, Aleksander Madry, Daniel A. Spielman, and Shang-Hua Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 273–282, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993636.1993674.
- 10 Michael B. Cohen, Brittany Terese Fasy, Gary L. Miller, Amir Nayyeri, Richard Peng, and Noel Walkington. 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*, SODA '14, pages 204–216, USA, 2014. Society for Industrial and Applied Mathematics.
- 11 Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. Solving sdd linear systems in nearly $m \log 1/2n$ time. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pages 343–352, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2591796.2591833.
- 12 Keenan Crane, Mathieu Desbrun Fernando de Goes, and Peter Schröder. Digital geometry processing with discrete exterior calculus. In *ACM SIGGRAPH 2013 courses*, SIGGRAPH '13, New York, NY, USA, 2013. ACM.
- 13 Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Persistent cohomology and circular coordinates. *Discret. Comput. Geom.*, 45(4):737–759, 2011. doi:10.1007/s00454-011-9344-x.
- 14 Tamal K. Dey. Computing height persistence and homology generators in \mathbb{R}^3 efficiently. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2649–2662. SIAM, 2019. doi:10.1137/1.9781611975482.164.
- 15 Ming Ding, Maximilian Probst Gutenberg, Rasmus Kyng, and Peng Zhang. Hardness Results for Laplacians of Simplicial Complexes via Sparse-Linear Equation Complete Gadgets. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- 16 Joel Friedman. Computing betti numbers via combinatorial laplacians. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 386–391, New York, NY, USA, 1996. ACM. doi:10.1145/237814.237985.

23:16 Hodge Decomposition and General Laplacian Solvers

- 17 Allen Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2000. URL: <https://cds.cern.ch/record/478079>.
- 18 Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, USA, 2nd edition, 2012.
- 19 Arun Jambulapati and Aaron Sidford. Ultrasparse ultrasparsifiers and faster laplacian system solvers. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, pages 540–559, USA, 2021.
- 20 Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Math. Program.*, 127(1):203–244, March 2011. doi:10.1007/s10107-010-0419-x.
- 21 Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. 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*, STOC '13, pages 911–920, New York, NY, USA, 2013. ACM. doi:10.1145/2488608.2488724.
- 22 Ioannis Koutis and Richard Miller, Gary L. and Peng. Approaching optimality for solving SDD linear systems. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 235–244, Washington, DC, USA, 2010. IEEE Computer Society. doi:10.1109/FOCS.2010.29.
- 23 Ioannis Koutis and Richard Miller, Gary L. and Peng. A nearly- $m \log n$ time solver for SDD linear systems. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 590–598, Washington, DC, USA, 2011. IEEE Computer Society. doi:10.1109/FOCS.2011.85.
- 24 Rasmus Kyng and Peng Zhang. Hardness results for structured linear systems. *SIAM J. Comput.*, 49(4), 2020. doi:10.1137/17M1161774.
- 25 Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: analysis and an algorithm. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2002. URL: <https://proceedings.neurips.cc/paper/2001/file/801272ee79cfde7fa5960571fee36b9b-Paper.pdf>.
- 26 Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., USA, 1998.
- 27 Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. doi:10.1109/34.868688.
- 28 Daniel Spielman. Spectral and algebraic graph theory. Available at <http://cs-www.cs.yale.edu/homes/spielman/sagt/sagt.pdf> (2021/12/01).
- 29 Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 563–568, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1374376.1374456.
- 30 Shang-Hua Spielman, Daniel A. and Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 81–90. ACM, 2004.
- 31 John Stillwell. *Classical Topology and Combinatorial Group Theory*, volume 72 of *Graduate Texts in Mathematics*. Springer, second edition, 1993. URL: <http://www.springer.com/mathematics/algebra/book/978-0-387-97970-0>.
- 32 Alireza Tahbaz-Salehi and Ali Jadbabaie. Distributed coverage verification in sensor networks without location information. *IEEE Transactions on Automatic Control*, 55(8):1837–1849, 2010. doi:10.1109/TAC.2010.2047541.
- 33 Yiyong Tong, Santiago Lombeyda, Anil N. Hirani, and Mathieu Desbrun. Discrete multiscale vector field decomposition. *ACM Trans. Graph.*, 22(3):445–452, July 2003. doi:10.1145/882262.882290.

- 34 Pravin M. Vaidya. Solving linear equations with symmetric diagonally dominant matrices by constructing good preconditioners. Workshop Talk at the IMA Workshop on Graph Theory and Sparse Matrix Computation, October 1991. Minneapolis, MN.
- 35 Qianqian Xu, Qingming Huang, Tingting Jiang, Bowei Yan, Weisi Lin, and Yuan Yao. Hodgerank on random graphs for subjective video quality assessment. *IEEE Transactions on Multimedia*, 14(3):844–857, 2012. doi:10.1109/TMM.2012.2190924.