



Quantifying the Structural Stability of Simplicial Homology

Nicola Guglielmi¹ · Anton Savostianov¹ · Francesco Tudisco¹

Received: 25 April 2023 / Revised: 24 July 2023 / Accepted: 26 July 2023 /

Published online: 14 August 2023

© The Author(s) 2023

Abstract

Simplicial complexes are generalizations of classical graphs. Their homology groups are widely used to characterize the structure and the topology of data in e.g. chemistry, neuroscience, and transportation networks. In this work we assume we are given a simplicial complex and that we can act on its underlying graph, formed by the set of 1-simplices, and we investigate the stability of its homology with respect to perturbations of the edges in such graph. Precisely, exploiting the isomorphism between the homology groups and the higher-order Laplacian operators, we propose a numerical method to compute the smallest graph perturbation sufficient to change the dimension of the simplex's Hodge homology. Our approach is based on a matrix nearness problem formulated as a matrix differential equation, which requires an appropriate weighting and normalizing procedure for the boundary operators acting on the Hodge algebra's homology groups. We develop a bilevel optimization procedure suitable for the formulated matrix nearness problem and illustrate the method's performance on a variety of synthetic quasi-triangulation datasets and real-world transportation networks.

Keywords Simplicial complexes · Homology groups · Hodge Laplacian · Matrix nearness problems · Matrix ODEs · Spectral optimization

Mathematics Subject Classification 05C50 · 65F45 · 65K10 · 57M15 · 62R40

1 Introduction

Models based on graphs are ubiquitous in the sciences and engineering as they allow us to model various complex systems in a unified form. Despite being widely and successfully used, graph-based models are limited to pairwise relationships and a variety of recent work

✉ Anton Savostianov
anton.savostianov@gssi.it

Nicola Guglielmi
nicola.guglielmi@gssi.it

Francesco Tudisco
francesco.tudisco@gssi.it

¹ Gran Sasso Science Institute, L'Aquila, Italy

has shown that many complex systems and datasets are better described by higher-order relations that go beyond pairwise interactions [5, 7, 9]. Relational data is full of interactions that happen in groups. For example, friendship relations often involve groups that are larger than two individuals and triangles are important building blocks of relational data [1, 4, 24]. Also in the presence of point-cloud data, directly modeling higher-order data interactions has led to improvements in numerous data mining settings, including clustering [7, 16, 34], link prediction [4, 8], and ranking [6, 32, 33].

A popular extension of dyadic models to the higher-order setting uses hypergraphs and relies on the transition from matrix-based to tensor-based models [5, 9]. Based on recent generalizations of graph spectral theory to tensor eigenproblems [17, 28], a number of graph methods based on matrices for e.g. clustering, community detection, centrality, opinion dynamics, have been extended to higher-order models based on tensors [4, 6, 25, 26, 34].

Alongside hypergraphs and tensors, *simplicial complexes* and *higher-order Laplacians* are another standard model for higher-order interactions, where simplices of different order can connect an arbitrary number of nodes [5, 9]. Higher-order Laplacians are linear operators that generalize the better-known graph Laplacian and provide key algebraic tools that allow to describe a simplicial complex and its structural properties. In particular, their kernels define a homology of the data and reveal fundamental topological properties such as connected components, holes, and voids [11, 22].

In this work we are interested to quantifying the stability of such topological properties, with respect to edge perturbations. More precisely, given an initial simplicial complex $\mathcal{K}$ with a corresponding homology and an underlying graph $\mathcal{G}_{\mathcal{K}}$ formed by the 1-simplices of $\mathcal{K}$, we want to quantify how far is $\mathcal{K}$ from another simplicial complex with a homology group of strictly different dimension. Here we assume we can act on the edges of $\mathcal{G}_{\mathcal{K}}$ and “being far” is measured in terms of the least number (or, in the weighted case, the least weight) of edges of $\mathcal{G}_{\mathcal{K}}$ that can be eliminated to change the dimension of the chosen homology group. While this form of stability is reminiscent of the persistence of the homology, which is widely studied in the literature, we remark that the two are significantly different. Rather than starting with a dataset of points on which to build a chain of simplices and an associated persistency diagram, in our setting, we assume we are given an initial simplex as a result of a data assimilation and modeling process we have no access to. Thus, acting on the elementary pairwise information we are given, the 1-simplices, we want to quantify the robustness of the homology that characterizes the simplex at hand. While a great effort has been devoted in recent years to measure the presence and persistence of simplicial homology [13, 27], much less is available about the stability of the homology classes with respect to data perturbation in the non-regular (non-geometric) setting [10].

The resulting problem can be formulated as a structured matrix nearness problem, which we approach by means of a spectral objective function and the integration of the corresponding matrix-valued gradient flow. In order to make a sound mathematical formulation of the problem we aim to solve, and of the numerical model we design for its solution, we structure the remainder of the paper as follows: in Sect. 2 we review in detail the notion of simplicial complexes and the corresponding higher-order Laplacians. In Sects. 2.2 and 3.1 we discuss how these operators may be extended to account for weighted higher-order node relations and formulate the corresponding stability problem in Sects. 3 and 4. After introducing a suitable spectral functional whose minimization to zero mathematically translates the stability problem, in Sects. 5 and 6 we present a two-level methodology whose inner iteration consists of a constrained matrix gradient flow on which is based our numerical method, and whose outer level tunes a suitable scalar parameter in order to solve a related scalar equation. Finally,

we devote Sect. 7 to illustrate the performance of the proposed numerical scheme on several example datasets.

2 Simplicial Complexes and Higher-Order Relations

A graph $\mathcal{G}$ is a pair of sets $(\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, n\}$ is the set of vertices and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is a set of unordered pairs representing the undirected edges of $\mathcal{G}$. We let m denote the number of edges $\mathcal{E} = \{e_1, \dots, e_m\}$ and we assume them ordered lexicographically, with the convention that $i < j$ for any $\{i, j\} \in \mathcal{E}$. For brevity, we often write ij in place of $\{i, j\}$ to denote the edge joining i and j . Moreover, we assume no self-loops, i.e., $ii \notin \mathcal{E}$ for all $i \in \mathcal{V}$.

A graph only considers pairwise relations between the vertices. A simplicial complex $\mathcal{K}$ is a generalization of a graph that allows us to model connections involving an arbitrary number of nodes by means of higher-order simplices. Formally, a k -th order simplex (or k -simplex, briefly) is a set of $k + 1$ vertices $\{i_0, i_1, \dots, i_k\}$ with the property that every subset of k nodes itself is a $(k - 1)$ -simplex. Any $(k - 1)$ -simplex of a k -simplex is called a *face*. The collection of all such simplices forms a simplicial complex $\mathcal{K}$, which therefore essentially consists of a collection of sets of vertices such that every subset of the set in the collection is in the collection itself. Thus, a graph $\mathcal{G}$ can be thought of as the collection of 0- and 1- simplices: the 0-simplices form the nodes set of $\mathcal{G}$, while 1-simplices form its edges. To emphasize this analogy, in the sequel we often specify that $\mathcal{K}$ is a simplicial complex on the vertex set $\mathcal{V}$.

Just like the edges of a graph, to any simplicial complex, we can associate an orientation (or ordering). To underline that an ordering has been fixed, we denote an ordered k -simplex σ using square brackets $\sigma = [i_0 \dots i_k]$. In particular, as for the case of edges, we always assume the lexicographical ordering, unless specified otherwise. That is, we assume that:

1. any k -simplex $[i_0 \dots i_k]$ in $\mathcal{K}$ is such that $i_0 < \dots < i_k$;
2. the k -simplices $\sigma_1, \sigma_2, \dots$ of $\mathcal{K}$ are ordered so that $\sigma_i < \sigma_{i+1}$ for all i , where $[i_0 \dots i_k] < [i'_0 \dots i'_k]$ if and only if there exists h such that $0 \leq h \leq k$, $i_j = i'_j$ for $j = 0, \dots, h$ and $i_h < i'_h$.

As for the edges, we often write $i_0 \dots i_k$ in place of $[i_0 \dots i_k]$ in this case.

2.1 Homology, Boundary Operators and Higher-Order Laplacians

Topological properties of a simplicial complex can be studied by considering boundary operators, higher-order Laplacians, and the associated homology. Here we recall these concepts trying to emphasize their matrix-theoretic flavor. To this end, we first fix some further notation and recall the notion of a real k -chain.

Definition 1 Assume $\mathcal{K}$ is a simplicial complex on the vertex set $\mathcal{V}$. For $k \geq 0$, we denote the set of all the oriented k -th order simplices in $\mathcal{K}$ as $\mathcal{V}_k(\mathcal{K})$ or simply $\mathcal{V}_k$. Thus, $\mathcal{V}_0 = \mathcal{V}$ and $\mathcal{V}_1 = \mathcal{E}$ form the underlying graph of $\mathcal{K}$, which we denote by $\mathcal{G}_{\mathcal{K}} = (\mathcal{V}, \mathcal{E}) = (\mathcal{V}_0, \mathcal{V}_1)$.

Definition 2 The formal real vector space spanned by all the elements of $\mathcal{V}_k$ with real coefficients is denoted by $C_k(\mathcal{K})$. Any element of $C_k(\mathcal{K})$, the formal linear combinations of simplices in $\mathcal{V}_k$, is called a k -chain.

We remark that, in the graph-theoretic terminology, $C_0(\mathcal{K})$ is usually called the space of *vertices' states*, while $C_1(\mathcal{K})$ is usually called the space of *flows* in the graph.




Fig. 1 Left-hand side panel: example of simplicial complex $\mathcal{K}$ on 7 nodes, and of the action of ∂_2 on the 2-simplex $[1, 2, 3]$; 2-simplices included in the complex are shown in red, arrows correspond to the orientation. Panels on the right: matrix forms B_1 and B_2 of boundary operators ∂_1 and ∂_2 respectively

The chain spaces are finite vector spaces generated by the set of k -simplices. The boundary and co-boundary operators are particular linear mappings between C_k and C_{k-1} , which in a way are the discrete analogous of high-order differential operators (and their adjoints) on continuous manifolds. The boundary operator ∂_k maps a k -simplex to an alternating sum of its $(k-1)$ -dimensional faces obtained by omitting one vertex at a time. Its precise definition is recalled below, while Fig. 1 provides an illustrative example of its action.

Definition 3 Let $k \geq 1$. Given a simplicial complex $\mathcal{K}$ over the set $\mathcal{V}$, the *boundary operator* $\partial_k : C_k(\mathcal{K}) \mapsto C_{k-1}(\mathcal{K})$ maps every ordered k -simplex $[i_0 \dots i_k] \in C_k(\mathcal{K})$ to the following alternated sum of its faces:

$$\partial_k[i_0 \dots i_k] = \sum_{j=0}^k (-1)^j [i_0 \dots i_{j-1} i_{j+1} \dots i_k] \in C_{k-1}(\mathcal{K})$$

As we assume the k -simplices in $\mathcal{V}_k$ are ordered lexicographically, we can fix a canonical basis for $C_k(\mathcal{K})$ and we can represent ∂_k as a matrix B_k with respect to such basis. In fact, once the ordering is fixed, $C_k(\mathcal{K})$ is isomorphic to $\mathbb{R}^{\mathcal{V}_k}$, the space of functions from $\mathcal{V}_k$ to $\mathbb{R}$ or, equivalently, the space of real vectors with $|\mathcal{V}_k|$ entries. Thus, B_k is a $|\mathcal{V}_{k-1}| \times |\mathcal{V}_k|$ matrix and ∂_k^* coincides with $B_k^\top$. We shall always assume the canonical basis for $C_k(\mathcal{K})$ is fixed in this way and we will deal exclusively with the matrix representation B_k from now on. An example of B_k for $k=1$ and $k=2$ is shown in Fig. 1.

A direct computation shows that the following fundamental identity holds (see e.g. [22, Thm. 5.7])

$$B_k B_{k+1} = 0 \quad (1)$$

for any k . This identity is also known in the continuous case as the Fundamental Lemma of Homology and it allows us to define a homology group associated with each k -chain. In fact, (1) implies in particular that $\text{im } B_{k+1} \subset \ker B_k$, so that the k -th homology group is correctly defined:

$$\mathcal{H}_k = \ker B_k / \text{im } B_{k+1}$$

The dimensionality of the k -th homology group is known as k -th Betti's number $\beta_k = \dim \mathcal{H}_k$, while the elements of $\mathcal{H}_k$ correspond to so-called k -dimensional holes in the simplicial complex. For example, $\mathcal{H}_0$, $\mathcal{H}_1$ and $\mathcal{H}_2$ describe *connected components*, *holes* and *three-dimensional voids* respectively. By standard algebraic passages one sees that $\mathcal{H}_k$ is isomorphic to $\ker (B_k^\top B_k + B_{k+1} B_{k+1}^\top)$. Thus, the number of k -dimensional holes corresponds to the dimension of the kernel of a linear operator, which is known as k -th order Laplacian or *higher-order Laplacian* of the simplicial complex $\mathcal{K}$.

Definition 4 Given a simplicial complex $\mathcal{K}$ and the boundary operators B_k and B_{k+1} , the k -th order Laplacian L_k of $\mathcal{K}$ is the $|\mathcal{V}_k| \times |\mathcal{V}_k|$ matrix defined as:

$$L_k = B_k^\top B_k + B_{k+1} B_{k+1}^\top \quad (2)$$

In particular, we remark that:

- The 0-order Laplacian is the standard *combinatorial graph Laplacian* $L_0 = B_1 B_1^\top \in \mathbb{R}^{n \times n}$, whose diagonal entries consist of the degrees of the corresponding vertices (i.e. the number of 1-simplices each vertex belongs to), while the off-diagonal $(L_0)_{ij}$ is equal to -1 if either ij is a 1-simplex, and it is zero otherwise;
- The 1-order Laplacian is known as *Hodge Laplacian* $L_1 = B_1^\top B_1 + B_2 B_2^\top \in \mathbb{R}^{m \times m}$. Similarly to the 0-order case L_0 , one can describe the entries of L_1 in terms of the structure of the simplicial complex, see e.g. [23].

2.1.1 Connected Components and Holes

The boundary operators B_k on $\mathcal{K}$ are directly connected with discrete notions of differential operators on the graph. In particular, B_1 , $B_1^\top$, and $B_2^\top$ are the graph's divergence, gradient, and curl operators, respectively. We refer to [22] for more details. As $\mathcal{H}_k$ is isomorphic to $\ker L_k$, we have that the following *Hodge decomposition* of $\mathbb{R}^{\mathcal{V}_k}$ holds

$$\mathbb{R}^{\mathcal{V}_k} = \text{im } B_{k+1} \oplus \text{im } B_k^\top \oplus \ker L_k = \text{im } \bar{B}_{k+1} \oplus \text{im } \bar{B}_k^\top \oplus \ker \bar{L}_k. \quad (3)$$

Thus, the space of vertex states $\mathbb{R}^{\mathcal{V}_0}$ can be decomposed as $\mathbb{R}^{\mathcal{V}_0} = \text{im } B_1 \oplus \ker L_0$, the sum of divergence-free vectors and harmonic vectors, which correspond to the connected components of the graph. In particular, for a connected graph, $\ker L_0$ is one-dimensional and consists of entry-wise constant vectors. Thus, for a connected graph, $\text{im } B_1$ is the set of vectors whose entries sum up to zero.

Similarly, the space of flows on graph's edges $\mathbb{R}^{\mathcal{V}_1}$ can be decomposed as $\mathbb{R}^{\mathcal{V}_1} = \text{im } B_2 \oplus \text{im } B_1^\top \oplus \ker L_1$. Thus, each flow can be decomposed into its gradient part $\text{im } B_1^\top$, which consists of flows with zero cycle sum, its curl part $\text{im } B_2$, which consists of circulations around order-2 simplices in $\mathcal{K}$, and its harmonic part $\ker L_1$, which represents 1-dimensional holes defined as global circulations modulo the curl flows.

While 0-dimensional holes are easily understood as the connected components of the graph $\mathcal{G}_\mathcal{K}$, a notion of "holes in the graph" $\mathcal{G}_\mathcal{K}$ corresponds to 1-dimensional holes in $\mathcal{K}$. This terminology comes from the analogy with the continuous case. In fact, if the graph is obtained as a discretization of a continuous manifold, harmonic functions in the homology group $\mathcal{H}_1$ would correspond to the holes in the manifold, as illustrated in Fig. 2. Moreover, the Hodge Laplacian of a simplicial complex built on N randomly sampled points in the manifold converges in the thermodynamic limit to its continuous counterpart, as $N \rightarrow \infty$, [12].




Fig. 2 Continuous and analogous discrete manifolds with one 1-dimensional hole ($\dim \overline{\mathcal{H}}_1 = 1$). Left pane: the continuous manifold; center pane: the discretization with mesh vertices; right pane: a simplicial complex built upon the mesh. Triangles in the simplicial complex $\mathcal{K}$ are colored gray (right) (Color figure online)

2.2 Normalized and Weighted Higher-Order Laplacians

In the classical graph setting, a normalized and weighted version of the Laplacian matrix is very often employed in applications. From a matrix theoretic point of view, having a weighted graph corresponds to allowing arbitrary nonnegative entries in the adjacency matrix defining the graph. In terms of boundary operators, this coincides with a positive diagonal rescaling. Analogously, the normalized Laplacian is defined by applying a diagonal congruence transformation to the standard Laplacian using the node weights. We briefly review these two constructions below.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with positive node and edge weight functions $w_0 : \mathcal{V} \rightarrow \mathbb{R}_{++}$ and $w_1 : \mathcal{E} \rightarrow \mathbb{R}_{++}$, respectively. Define the $|\mathcal{V}| \times |\mathcal{V}|$ and $|\mathcal{E}| \times |\mathcal{E}|$ diagonal matrices $W_0 = \text{Diag}\{w_0(v_i)^{1/2}\}_{i=1}^n$ and $W_1 = \text{Diag}\{w_1(e_i)^{1/2}\}_{i=1}^m$. Then $\overline{B}_1 = W_0^{-1} B_1 W_1$ is the normalized and weighted boundary operator of $\mathcal{G}$ and we have that

$$\overline{L}_0 = \overline{B}_1 \overline{B}_1^\top = W_0^{-1} B_1 W_1^2 B_1^\top W_0^{-1} \quad (4)$$

is the normalized weighted graph Laplacian of $\mathcal{G}$.

In particular, note that, as for L_0 , the entries of $\overline{L}_0$ uniquely characterize the graph topology, in fact we have

$$(\overline{L}_0)_{ii} = \frac{\deg(i)}{w_0(i)}, \quad (\overline{L}_0)_{ij} = \begin{cases} -\frac{w_1(ij)}{\sqrt{w_0(i)w_0(j)}} & ij \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}, \quad \text{for } i \neq j$$

where $\deg(i)$ denotes the (weighted) degree of node i , i.e. $\deg(i) = \sum_{e \in \mathcal{E}: i \in e} w_1(e)$.

While the definition of k -th order Laplacian is well-established for the case of unweighted edges and simplices, a notion of weighted and normalized k -th order Laplacian is not universally available and it might depend on the application one has at hand. For example, different definitions of weighted Hodge Laplacian are considered in [11, 20, 22, 29].

At the same time, we notice that the notation used in (4) directly generalizes to higher orders. Thus, we propose the following notion of normalized and weighted k -th Laplacian

Definition 5 Let $\mathcal{K}$ be a simplicial complex and let $w_k : \mathcal{V}_k \rightarrow \mathbb{R}_{++}$ be a positive-valued weight function on the k -simplices of $\mathcal{K}$. Define the diagonal matrix $W_k = \text{Diag}\{w_k(\sigma_i)^{1/2}\}_{i=1}^{|\mathcal{V}_k|}$. Then, $\overline{B}_k = W_{k-1}^{-1} B_k W_k$ is the normalized and weighted k -th boundary operator, to which corresponds the normalized and weighted k -th Laplacian

$$\begin{aligned} \overline{L}_k &= \overline{B}_k^\top \overline{B}_k + \overline{B}_{k+1} \overline{B}_{k+1}^\top \\ &= W_k B_k^\top W_{k-1}^{-2} B_k W_k + W_k^{-1} B_{k+1} W_{k+1}^2 B_{k+1}^\top W_k^{-1}. \end{aligned} \quad (5)$$

Note that, from the definition $\overline{B}_k = W_{k-1}^{-1} B_k W_k$ and (1), we immediately have that $\overline{B}_k \overline{B}_{k+1} = 0$. Thus, the group $\overline{\mathcal{H}}_k = \ker \overline{B}_k / \text{im } \overline{B}_{k+1}$ is well defined for any choice of

positive weights w_k and is isomorphic to $\ker \bar{L}_k$. While the homology group may depend on the weights, we observe below that its dimension does not. Precisely, we have

Proposition 1 *The dimension of the homology groups of $\mathcal{K}$ is not affected by the weights of its k -simplices. Precisely, if W_k are positive diagonal matrices, we have*

$$\dim \ker \bar{B}_k = \dim \ker B_k, \quad \dim \ker \bar{B}_k^\top = \dim \ker B_k^\top, \quad \dim \bar{\mathcal{H}}_k = \dim \mathcal{H}_k. \quad (6)$$

Moreover, $\ker B_k = W_k \ker \bar{B}_k$ and $\ker B_k^\top = W_{k-1}^{-1} \ker \bar{B}_k^\top$.

Proof Since W_k is an invertible diagonal matrix,

$$\bar{B}_k \mathbf{x} = 0 \iff W_{k-1}^{-1} B_k W_k \mathbf{x} = 0 \iff B_k W_k \mathbf{x} = 0.$$

Hence, if $\mathbf{x} \in \ker \bar{B}_k$, then $W_k \mathbf{x} \in \ker B_k$, and, since W_k is bijective, $\dim \ker \bar{B}_k = \dim \ker B_k$. Similarly, one observes that $\dim \ker \bar{B}_k^\top = \dim \ker B_k^\top$. Moreover, since $\bar{B}_k \bar{B}_{k+1} = 0$, then $\text{im } \bar{B}_{k+1} \subseteq \ker \bar{B}_k$ and $\text{im } \bar{B}_k^\top \subseteq \ker \bar{B}_{k+1}^\top$. This yields $\ker \bar{B}_k \cup \ker \bar{B}_{k+1}^\top = \mathbb{R}^{\mathcal{V}_k} = \ker B_k \cup \ker B_{k+1}^\top$. Thus, for the homology group it holds $\dim \bar{\mathcal{H}}_k = \dim \mathcal{H}_k$. $\square$

2.3 Principal Spectral Inheritance

Before moving on to the next section, we recall here a relatively direct but important spectral property that connects the spectra of the k -th and $(k+1)$ -th order Laplacians.

Theorem 1 (Spectral inheritance of higher-order Laplacians) *Let $\bar{L}_k^{\text{down}} = \bar{B}_k^\top \bar{B}_k$ and $\bar{L}_k^{\text{up}} = \bar{B}_{k+1} \bar{B}_{k+1}^\top$. Then:*

1. $\sigma_+(\bar{L}_k^{\text{up}}) = \sigma_+(\bar{L}_{k+1}^{\text{down}})$, where $\sigma_+(\cdot)$ denotes the set of positive eigenvalues;
2. for any $\mu \in \sigma_+(\bar{L}_k)$, either $\mu \in \sigma_+(\bar{L}_k^{\text{up}})$ or the corresponding eigenvector $\mathbf{v} \in \ker \bar{L}_k^{\text{up}}$. Similarly, for any $\nu \in \sigma_+(\bar{L}_{k+1})$, either $\nu \in \sigma_+(\bar{L}_{k+1}^{\text{down}})$ or the corresponding eigenvector $\mathbf{u} \in \ker \bar{L}_{k+1}^{\text{down}}$, and

$$\bar{B}_k^\top \bar{B}_k \mathbf{v} = \mu \mathbf{v}, \quad \bar{B}_{k+2} \bar{B}_{k+2}^\top \mathbf{u} = \nu \mathbf{u}.$$

Proof For $\mu > 0$ it is sufficient to note that if $(\mu, \mathbf{x})$ is an eigenpair for $\bar{L}_k^{\text{up}}$, then $(\mu, \mu^{-1/2} \bar{B}_{k+1}^\top \mathbf{x})$ is an eigenpair for $\bar{L}_{k+1}^{\text{down}}$. Similarly, if $(\mu, \mathbf{y})$ is an eigenpair of $\bar{L}_{k+1}^{\text{down}}$, then $(\mu, \mu^{-1/2} \bar{B}_{k+1} \mathbf{y})$ is the corresponding eigenpair of $\bar{L}_k^{\text{up}}$, yielding (1). The statement in (2) follows immediately from the Hodge decomposition (3). $\square$

In other words, the variation of the spectrum of the k -th Laplacian when moving from one order to the next one works as follows: the down-term $\bar{L}_{k+1}^{\text{down}}$ inherits the positive part of the spectrum from the up-term of $\bar{L}_k^{\text{up}}$; the eigenvectors corresponding to the inherited positive part of the spectrum lie in the kernel of $\bar{L}_{k+1}^{\text{up}}$; at the same time, the “new” up-term $\bar{L}_{k+1}^{\text{up}}$ has a new, non-inherited, part of the positive spectrum (which, in turn, lies in the kernel of the $(k+2)$ -th down-term).

In particular, we notice that for $k=0$, since $B_0=0$ and $\bar{L}_0=\bar{L}_0^{\text{up}}$, the theorem yields $\sigma_+(\bar{L}_0)=\sigma_+(\bar{L}_1^{\text{down}})\subseteq\sigma_+(\bar{L}_1)$. In other terms, the positive spectrum of the $\bar{L}_0$ is inherited by the spectrum of $\bar{L}_1$ and the remaining (non-inherited) part of $\sigma_+(\bar{L}_1)$ coincides with $\sigma_+(\bar{L}_1^{\text{up}})$. Figure 3 provides an illustration of the statement of Theorem 1 for $k=0$.




Fig. 3 Illustration for the principal spectrum inheritance (Theorem 1) in case $k = 0$: spectra of $\bar{L}_1$, $\bar{L}_1^{down}$ and $\bar{L}_1^{up}$ are shown. Colors signify the splitting of the spectrum, $\lambda_i > 0 \in \sigma(\bar{L}_1)$; all yellow eigenvalues are inherited from $\sigma_+(\bar{L}_0)$; red eigenvalues belong to the non-inherited part. Dashed barrier μ signifies the penalization threshold (see the target functional in Sect. 4) preventing homological pollution (see Sect. 3.1) (Color figure online)

3 Problem Setting: Nearest Complex with Different Homology

Suppose we are given a simplicial complex $\mathcal{K}$ on the vertex set $\mathcal{V}$, with simplex weight functions $w_0, w_1, \dots$, and let $\beta_k = \dim \mathcal{H}_k = \dim \bar{\mathcal{H}}_k$ the dimension of its k -homology. We aim at finding the closest simplex on the same vertex set $\mathcal{V}$, with a strictly larger number of holes. As it is the most frequent higher-order Laplacian appearing in applications and since this provides already a large number of numerical complications, we focus here only on the Hodge Laplacian case: given the simplicial complex $\mathcal{K} = (\mathcal{V}_0, \mathcal{V}_1, \mathcal{V}_2, \dots)$, we look for the smallest perturbation of the edges $\mathcal{V}_1$ which increases the dimension of $\mathcal{H}_1$. More precisely:

Problem 1 Let $\mathcal{K}$ be a simplicial complex of order at least 2 with associated edge weight function w_1 and corresponding diagonal weight matrix W_1 , and let $\beta_1(W_1)$ be the dimension of the homology group corresponding to the weights in W_1 . For $\varepsilon > 0$, let

$$\Omega(\varepsilon) = \left\{ \text{diagonal matrices } W \text{ such that } \|W\| = \varepsilon \right\},$$

$$\Pi(W_1) = \left\{ \text{diagonal matrices } W \text{ such that } W_1 + W \geq 0 \right\}.$$

In other words, $\Omega(\varepsilon)$ is an ε -sphere and $\Pi(W_1)$ allows only non-negative simplex weights. We look for the smallest perturbation ε such that there exists a weight modification $\delta W_1 \in \Omega(\varepsilon) \cap \Pi(W_1)$ such that $\beta_1(W_1) < \beta_1(W_1 + \delta W_1)$.

Here, and throughout the paper, $\|X\|$ denotes either the Frobenius norm if X is a matrix, or the Euclidean norm if X is a vector. Note that, as we are looking for the smallest ε , the equality $\|W\| = \varepsilon$ is an obvious choice, as opposed to $\|W\| \leq \varepsilon$.

As the dimension β_1 coincides with the dimension of the kernel of $\bar{L}_1$, we approach this problem through the minimization of a functional based on the spectrum of the 0-th and 1-st Laplacian of the simplicial complex. In order to define such functional, we first make a number of considerations.

Note that, due to Theorem 1, the dimension of the first homology group does not change when the edge weights are perturbed, as long as all the weights remain positive. Thus, in order to find the desired perturbation δW_1 , we need to set some of the initial weights to zero. This operation creates several potential issues we need to address, as discussed next.

First, setting an edge to zero implies that one is formally removing that edge from the simplicial complex. As the simplicial complex structure needs to be maintained, when doing so we need to set to zero also the weight of any 2-simplex that contains that edge. For this

reason, if $\tilde{w}_1(e) = w_1(e) + \delta w_1(e)$ is the new edge weight function, we require the weight function of the 2-simplices to change into $\tilde{w}_2$, defined as

$$\tilde{w}_2(i_1 i_2 i_3) = f\left(\frac{\delta w_1(i_1 i_2)}{w_1(i_1 i_2)}, \frac{\delta w_1(i_2 i_3)}{w_1(i_2 i_3)}, \frac{\delta w_1(i_1 i_3)}{w_1(i_1 i_3)}\right) \cdot w_2(i_1 i_2 i_3)$$

where $f(u_1, u_2, u_3)$ is a function such that $f(0, 0, 0) = 1$ and that monotonically decreases to zero as $u_i \rightarrow -1$, for any $i = 1, 2, 3$. An example of such f is

$$f(u_1, u_2, u_3) = 1 - \min\{u_1, u_2, u_3\}. \quad (7)$$

Second, when setting the weight of an edge to zero we may disconnect the underlying graph and create an all-zero column and row in the Hodge Laplacian. This gives rise to the phenomenon that we call “homological pollution”, which we will discuss in detail in the next subsection.

3.1 Homological Pollution: Inherited Almost Disconnectedness

As the dimension of Hodge homology β_1 corresponds to the number of zero eigenvalues in $\overline{L}_1$, the intuition suggests that if $\overline{L}_1$ has some eigenvalue that is close to zero, then the simplicial complex is “close to” having at least one more 1-dimensional hole. There are a number of problems with this intuitive consideration.

By Theorem 1 for $k = 0$, $\sigma_+(\overline{L}_1)$ inherits $\sigma_+(\overline{L}_0)$. Hence, if the weights in W_1 vary continuously so that a positive eigenvalue in $\sigma_+(\overline{L}_0)$ approaches 0, the same happens to $\sigma_+(\overline{L}_1)$. Assuming the initial graph $\mathcal{G}_K$ is connected, an eigenvalue that approaches zero in $\sigma(\overline{L}_0)$ would imply that $\mathcal{G}_K$ is approaching disconnectedness. This leads to a sort of *pollution* of the kernel of $\overline{L}_1$ as an almost-zero eigenvalue which corresponds to an “almost” 0-dimensional hole (disconnected component) from $\overline{L}_0$ is inherited into the spectrum of $\overline{L}_1$, but this small eigenvalue of $\overline{L}_1$ does not correspond to the creation of a new 1-dimensional hole in the reduced complex.

To better explain the problem of homological pollution, let us consider the following illustrative example.

Example 1 Consider the simplicial complex of order 2 depicted in Fig. 4a. In this example we have $\mathcal{V}_0 = \{1, \dots, 7\}$, $\mathcal{V}_1 = \{[1, 2], [1, 3], [2, 3], [2, 4], [3, 5], [4, 5], [4, 6], [5, 6], [5, 7], [6, 7]\}$ and $\mathcal{V}_2 = \{[1, 2, 3], [4, 5, 6], [5, 6, 7]\}$, all with weight equal to one: $w_k \equiv 1$ for




Fig. 4 Example of the homological pollution, Example 1, for the simplicial complex $\mathcal{K}$ on 7 vertices; the existing hole is $[2, 3, 4, 5]$ (left and center pane), all 3 cliques are included in the simplicial complex and shown in blue. The left pane demonstrates the initial setup with 1 hole; the center pane retains the hole exhibiting spectral pollution; the continuous transition to the eliminated edges with $\beta_1 = 0$ (no holes) is shown on the right pane

$k = 0, 1, 2$. The only existing 1-dimensional hole is shown in red and thus the corresponding Hodge homology is $\beta = 1$. Now, consider perturbing the weight of edges $[2, 4]$ and $[3, 5]$ by setting their weights to $\varepsilon > 0$ Fig. 4b. For small ε , this perturbation implies that the smallest nonzero eigenvalue μ_2 in $\sigma_+(\bar{L}_0)$ is scaled by ε . As $\sigma_+(\bar{L}_0) \subseteq \sigma_+(\bar{L}_1)$, we have that $\dim \ker \bar{L}_1 = 1$ and $\sigma_+(\bar{L}_1)$ has an arbitrary small eigenvalue, approaching 0 with $\varepsilon \rightarrow 0$. At the same time, when $\varepsilon \rightarrow 0$, the reduced complex obtained by removing the zero edges as in Fig. 4c does not have any 1-dimensional hole, i.e. $\beta_1 = 0$. Thus, in this case, the presence of a very small eigenvalue $\mu_2 \in \sigma_+(\bar{L}_1)$ does not imply that the simplicial complex is close to a simplicial complex with a larger dimension of the Hodge homology.

To mitigate the phenomenon of homological pollution, in our spectral-based functional for Problem 1 we include a term that penalizes the spectrum of $\bar{L}_0$ from approaching zero. To this end, we observe below that a careful choice of the vertex weights is required.

The smallest non-zero eigenvalue of the Laplacian $\mu_2 \in \sigma(\bar{L}_0)$ is directly related to the connectedness of the graph $\mathcal{G}_K$. This relation is well-known and dates back to the pioneering work of Fiedler [14]. In particular, as μ_2 is a function of node and edge weights, the following generalized version of the Cheeger inequality holds (see e.g. [31])

$$\frac{1}{2}\mu_2 \leq h(\mathcal{G}_K) \leq \left(2\mu_2 \max_{i \in \mathcal{V}_0} \frac{\deg(i)}{w_0(i)}\right)^{1/2}, \quad (8)$$

where $h(\mathcal{G}_K)$ is the Cheeger constant of the graph $\mathcal{G}_K$, defined as

$$h(\mathcal{G}_K) = \min_{S \subset \mathcal{V}_0} \frac{w_1(S, \mathcal{V}_0 \setminus S)}{\min\{w_0(S), w_0(\mathcal{V}_0 \setminus S)\}},$$

with $w_1(S, \mathcal{V}_0 \setminus S) = \sum_{i \in \mathcal{V}_1: i \in S, j \notin S} w_1(ij)$, $\deg(i) = \sum_{j: ij \in \mathcal{V}_1} w_1(ij)$, and $w_0(S) = \sum_{i \in S} w_0(i)$.

We immediately see from (8) that when the graph $\mathcal{G}_K$ is disconnected, then $h(\mathcal{G}_K) = 0$ and $\mu_2 = 0$ as well. Vice-versa, if μ_2 goes to zero, then $h(\mathcal{G}_K)$ decreases to zero too. While this happens independently of w_0 and w_1 , if w_0 is a function of w_1 then it might fail to capture the presence of edges whose weight is decreasing and is about to disconnect the graph.

To see this, consider the example choice $w_0(i) = \deg(i)$, the degree of node i in $\mathcal{G}_K$. Note that this is a very common choice in the graph literature, with several useful properties, including the fact that no other graph-dependent constant appears in the Cheeger inequality (8) other than μ_2 . For this weight choice, consider the case of a leaf node, a node $i \in \mathcal{V}_0$ that has only one edge $ij_0 \in \mathcal{V}_1$ connecting i to the rest of the (connected) graph $\mathcal{G}_K$ via the node j_0 . If we set $w_1(ij_0) = \varepsilon$ and we let ε decrease to zero, the graph $\mathcal{G}_K$ is approaching disconnectedness and we would expect $h(\mathcal{G}_K)$ and μ_2 to decrease as well. However, one easily verifies that both μ_2 and $h(\mathcal{G}_K)$ are constant with respect to ε in this case, as long as $\varepsilon \neq 0$.

In order to avoid such a discontinuity, in our weight perturbation strategy for the simplex $\mathcal{K}$, if w_0 is a function of w_1 , we perturb it by a constant shift. Precisely, if w_0 is the initial vertex weight of $\mathcal{K}$, we set $\tilde{w}_0(i) = w_0(i) + \varrho$, with $\varrho > 0$. So, for example, if $w_0 = \deg$ and the new edge weight function $\tilde{w}_1(e) = w_1(e) + \delta w_1(e)$ is formed after the addition of δW_1 , we set $\tilde{w}_0(i) = \sum_j [w_1(ij) + \delta w_1(ij)] + \varrho$.

4 Spectral Functional for 1-st Homological Stability

We are now in the position to formulate our proposed spectral-based functional, whose minimization leads to the desired small perturbation that changes the first homology of $\mathcal{K}$. In the notation of Problem 1, we are interested in the smallest perturbation ε and the corresponding modification $\delta W_1 \in \Omega(\varepsilon) \cap \Pi(W_1)$ that increases β_1 .

As $\|\delta W_1\| = \varepsilon$, for convenience we indicate $\delta W_1 = \varepsilon E$ with $\|E\| = 1$ so $E \in \Omega(1) \cap \Pi_\varepsilon(W_1)$, where $\Pi_\varepsilon(W_1) = \{W \mid \varepsilon W \in \Pi(W_1)\}$. For the sake of simplicity, we will omit the dependencies and write Ω and Π_ε , when there is no danger of ambiguity. Finally, let us denote by $\beta_1(\varepsilon, E)$ the first Betti number corresponding to the simplicial complex perturbed via the edge modification εE . With this notation, we can reformulate Problem 1 as follows:

Problem 2 Find the smallest $\varepsilon > 0$, such that there exists an admissible perturbation $E \in \Omega \cap \Pi_\varepsilon$ with an increased number of holes, i.e.

$$\min \{ \varepsilon > 0 : \beta_1(\varepsilon, E) \geq \beta_1 + 1 \text{ for some } E \in \Omega \cap \Pi_\varepsilon \} \quad (9)$$

where $\beta_1 = \beta_1(0, \cdot)$ is the first Betti number of the original simplicial complex.

In order to approach Problem 2, we introduce a target functional $F(\varepsilon, E)$, based on the spectrum of the 1-Laplacian $\bar{L}_1(\varepsilon, E)$ and the 0-Laplacian $\bar{L}_0(\varepsilon, E)$, where the dependence on ε and E is to emphasize the corresponding weight perturbation is of the form $W_1 \mapsto W_1 + \varepsilon E$.

Our goal is to move a positive entry from $\sigma_+(\bar{L}_1(\varepsilon, E))$ to the kernel. At the same time, assuming the initial graph $\mathcal{G}_{\mathcal{K}}$ is connected, one should avoid the inherited almost disconnectedness with small positive entries of $\sigma_+(\bar{L}_0(\varepsilon, E))$. As, by Theorem 1 for $k = 0$, $\sigma_+(\bar{L}_0(\varepsilon, E)) = \sigma_+(\bar{L}_1^{\text{down}}(\varepsilon, E))$, the only eigenvalue of $\bar{L}_1(\varepsilon, E)$ that can be continuously driven to 0 comes from $\bar{L}_1^{\text{up}}(\varepsilon, E)$. For this reason, let us denote the *first non-zero eigenvalue* of the up-Laplacian $\bar{L}_1^{\text{up}}(\varepsilon, E)$ by $\lambda_+(\varepsilon, E)$. The proposed target functional is defined as:

$$F(\varepsilon, E) = \frac{\lambda_+(\varepsilon, E)^2}{2} + \frac{\alpha}{2} \max \left(0, 1 - \frac{\mu_2(\varepsilon, E)}{\mu} \right)^2 \quad (10)$$

where α and μ are positive parameters, and $\mu_2(\varepsilon, E)$ is the first nonzero eigenvalue of $\bar{L}_0(\varepsilon, E)$. As recalled in Sect. 3.1, $\mu_2(\varepsilon, E)$ is an algebraic measure of the connectedness of the perturbed graph, thus the whole second term in (10) “activates” when such algebraic connectedness falls below the threshold μ .

By design, $F(\varepsilon, E)$ is non-negative and is equal to 0 iff $\lambda_+(\varepsilon, E)$ reaches 0, increasing the dimension of $\mathcal{H}_1$. Using this functional, we recast the Problem 2 as

$$\min \{ \varepsilon > 0 : F(\varepsilon, E) = 0 \text{ for some } E \in \Omega_\varepsilon \} \quad (11)$$

Remark 1 When $\mathcal{G}_{\mathcal{K}}$ is connected, $\dim \ker \bar{L}_0 = 1$ and, by the Theorem 1, $\dim \ker \bar{L}_1^{\text{up}} = \dim \ker \bar{L}_1 + (n - \dim \ker \bar{L}_0) = n + \beta_1 - 1$, so the first nonzero eigenvalue of $\bar{L}_1^{\text{up}}$ is the $(n + \beta_1)$ -th. While $(n + \beta_1)$ can be a large number in practice, we will discuss in Sect. 6.1 an efficient method that allows us to compute $\lambda_+(\varepsilon, E)$ without computing any of the previous $(n + \beta_1 - 1)$ eigenvalues.

5 A Two-Level Optimization Procedure

We approach problem (11) by the means of a two-level iterative method, which is based on the successive minimization of the target functional $F(\varepsilon, E)$ and a subsequent tuning of the parameter ε . More precisely, we propose the following two-level scheme.

1. *Inner level* for fixed $\varepsilon > 0$, solve the minimization problem

$$E(\varepsilon) = \arg \min_{E \in \Omega \cap \Pi_\varepsilon} F(\varepsilon, E)$$

by a constrained gradient flow which we formulate below, where we denote the computed minimizer by $E(\varepsilon)$.

2. *Outer level* given the function $\varepsilon \mapsto E(\varepsilon)$, we wish to solve the equation:

$$F(\varepsilon, E(\varepsilon)) = 0 \quad (12)$$

and our goal is to compute the smallest solution $\varepsilon^* > 0$ of (12).

A similar procedure was used in the context of graph spectral nearness in [2, 3] and in other matrix nearness problems [19].

We approach the inner level by means of the constrained gradient system

$$\dot{E}(t) = -\mathbb{P}_{\Omega \cap \Pi_\varepsilon} G(\varepsilon, E(t)), \quad G(\varepsilon, E) = \nabla_E F(\varepsilon, E) \quad (13)$$

where $\mathbb{P}_{\Omega \cap \Pi_\varepsilon}$ is an orthogonal projector (wrt to Frobenius inner product) onto the admissible set $\Omega \cap \Pi_\varepsilon$ (where ε is fixed). Since the system integrates the anti-gradient, a minimizer (at least local) of $F(\varepsilon, E)$ is obtained as $t \rightarrow \infty$.

We devote the next two Sect. 5.1 and Sect. 5.2 to computing the projected gradient in (13). The idea is to express the derivative of F in terms of the derivative of the perturbation and this way identifying the free gradient, and then projecting it onto the constraints obtaining the constrained gradient. To this end, we first compute the free gradient and then we discuss how to deal with the projection onto the admissible set.

By construction, the resulting algorithm converges to a (local) minimum of $F(\varepsilon, E)$. Although global convergence to the global optimum is not guaranteed, in a few simple experiments we observe the method reaches the expected global solution.

5.1 The Free Gradient

We compute here the free gradient of F with respect to E , given a fixed ε . In order to proceed, we need a few preliminary results.

The following is a standard perturbation result for eigenvalues; see e.g. [21], where we denote by $\langle X, Y \rangle = \sum_{i,j} x_{ij} y_{ij} = \text{Tr}(X^T Y)$ the inner product in $\mathbb{R}^{n \times n}$ that induces the Frobenius norm $\|X\| = \langle X, X \rangle^{1/2}$.

Theorem 2 (Derivative of simple eigenvalues) *Consider a continuously differentiable path of square symmetric matrices $A(t)$ for t in an open interval $\mathcal{I}$. Let $\lambda(t)$, $t \in \mathcal{I}$, be a continuous path of simple eigenvalues of $A(t)$. Let $\mathbf{x}(t)$ be the eigenvector associated to the eigenvalue $\lambda(t)$ and assume $\|\mathbf{x}(t)\| = 1$ for all t . Then λ is continuously differentiable on $\mathcal{I}$ with the derivative (denoted by a dot)*

$$\dot{\lambda} = \mathbf{x}^T \dot{A} \mathbf{x} = \langle \mathbf{x} \mathbf{x}^T, \dot{A} \rangle. \quad (14)$$

Moreover, “continuously differentiable” can be replaced with “analytic” in the assumption and the conclusion.

Let us denote the perturbed weight matrix by $\tilde{W}_1(t) = W_1 + \varepsilon E(t)$, and the corresponding $\tilde{W}_0(t) = W_0(\tilde{W}_1(t))$ and $\tilde{W}_2(t) = W_2(\tilde{W}_1(t))$, defined accordingly as discussed in Sect. 3. From now on we omit the time dependence for the perturbed matrices to simplify the notation. Since $\tilde{W}_0$, $\tilde{W}_1$ and $\tilde{W}_2$ are necessarily diagonal, by the chain rule we have $\dot{\tilde{W}}_i(t) = \varepsilon \text{diag}\left(J_1^i \dot{E} \mathbf{1}\right)$, where $\mathbf{1}$ is the vector of all ones, $\text{diag}(\mathbf{v})$ is the diagonal matrix with diagonal entries the vector $\mathbf{v}$, and J_1^i is the Jacobian matrix of the i -th weight matrix with respect to $\tilde{W}_1$, which for any $u_1 \in \mathcal{V}_1$ and $u_2 \in \mathcal{V}_i$, has entries $[J_1^i]_{u_1, u_2} = \frac{\partial}{\partial \tilde{w}_1(u_1)} \tilde{w}_i(u_2)$.

Next, in the following two lemmas, we express the time derivative of the Laplacian $\bar{L}_0$ and $\bar{L}_1^{up}$ as functions of $E(t)$. The proofs of these results are straightforward and omitted for brevity. In what follows, $\text{Sym}[A]$ denotes the symmetric part of the matrix A , namely $\text{Sym}[A] = (A + A^\top)/2$.

Lemma 1 (Derivative of $\bar{L}_0$) For the simplicial complex $\mathcal{K}$ with the initial edges’ weight matrix W_1 and fixed perturbation norm ε , let $E(t)$ be a smooth path and $\tilde{W}_0, \tilde{W}_1, \tilde{W}_2$ be corresponding perturbed weight matrices. Then,

$$\frac{1}{2\varepsilon} \frac{d}{dt} \bar{L}_0(t) = \tilde{W}_0^{-1} B_1 \tilde{W}_1 \dot{E} B_1^\top \tilde{W}_0^{-1} - \text{Sym} \left[\tilde{W}_0^{-1} \text{diag} \left(J_1^0 \dot{E} \mathbf{1} \right) \bar{L}_0 \right]. \quad (15)$$

Lemma 2 (Derivative of $\bar{L}_1^{up}$) For the simplicial complex $\mathcal{K}$ with the initial edges’ weight matrix W_1 and fixed perturbation norm ε , let $E(t)$ be a smooth path and $\tilde{W}_0, \tilde{W}_1, \tilde{W}_2$ be corresponding perturbed weight matrices. Then,

$$\frac{1}{2\varepsilon} \frac{d}{dt} \bar{L}_1^{up}(t) = -\text{Sym} \left[\tilde{W}_1^{-1} B_2 \tilde{W}_2^2 B_2^\top \tilde{W}_1^{-1} \dot{E} \tilde{W}_1^{-1} \right] + \tilde{W}_1^{-1} B_2 \tilde{W}_2 \text{diag} \left(J_1^0 \dot{E} \mathbf{1} \right) B_2^\top \tilde{W}_1^{-1}$$

Combining Lemma 2 with Lemma 1 and Lemma 2 we obtain the following expression for the free gradient of the functional.

Theorem 3 (The free gradient of $F(\varepsilon, E)$) Assume the initial weight matrices W_0, W_1 and W_2 , as well as the parameters $\varepsilon > 0$, $\alpha > 0$ and $\mu > 0$, are given. Additionally assume that $E(t)$ is a differentiable matrix-valued function such that the first non-zero eigenvalue $\lambda_+(\varepsilon, E)$ of $\bar{L}_1^{up}(\varepsilon, E)$ and the second smallest eigenvalue $\mu_2(\varepsilon, E)$ of $\bar{L}_0(\varepsilon, E)$ are simple. Let $\tilde{W}_0, \tilde{W}_1, \tilde{W}_2$ be corresponding perturbed weight matrices; then:

$$\begin{aligned} \frac{1}{\varepsilon} \nabla_E F(\varepsilon, E)(t) &= \lambda_+(\varepsilon, E) \cdot \\ &\cdot \left[\text{Sym} \left[-\tilde{W}_1^{-1} B_2 \tilde{W}_2^2 B_2^\top \tilde{W}_1^{-1} \mathbf{x}_+ \mathbf{x}_+^\top \tilde{W}_1^{-1} \right] \right. \\ &+ \text{diag} \left(J_1^{2^\top} \text{diagvec} \left(B_2^\top \tilde{W}_1^{-1} \mathbf{x}_+ \mathbf{x}_+^\top \tilde{W}_1^{-1} B_2 \tilde{W}_2 \right) \right) \left. - \frac{\alpha}{\mu} \max \left\{ 0, 1 - \frac{\mu_2(\varepsilon, E)}{\mu} \right\} \right] \cdot \\ &\cdot \left[B_1^\top \tilde{W}_0^{-1} \mathbf{y}_2 \mathbf{y}_2^\top \tilde{W}_0^{-1} B_1 \tilde{W}_1 - \text{diag} \left(J_1^{0^\top} \text{diagvec} \left(\text{Sym}[\tilde{W}_0^{-1} \mathbf{y}_2 \mathbf{y}_2^\top \bar{L}_0] \right) \right) \right] \end{aligned}$$

where $\mathbf{x}_+$ is a unit eigenvector of $\bar{L}_1^{up}$ corresponding to λ_+ , $\mathbf{y}_2$ is a unit eigenvector of $\bar{L}_0$ corresponding to μ_2 , and the operator $\text{diagvec}(X)$ returns the main diagonal of X as a vector.

Proof To derive the expression for the gradient $\nabla_E F$, we exploit the chain rule for the time derivative: $\dot{\lambda} = \langle \frac{d}{dt} A(E(t)), \mathbf{x} \mathbf{x}^\top \rangle = \langle \nabla_E \lambda, \dot{E} \rangle$. Then it is sufficient to apply the cyclic perturbation for the scalar products of Lemma 1 and Lemma 2 with $\mathbf{x}_+ \mathbf{x}_+^\top$ and $\mathbf{y}_2 \mathbf{y}_2^\top$ respectively. The final transition requires the formula:

$$\langle A, \text{diag}(BE\mathbf{1}) \rangle = \left\langle \text{diag} \left(B^\top (\text{diagvec } A) \right), E \right\rangle.$$

□

Remark 2 The derivation above assumes the simplicity of both $\mu_2(\varepsilon, E)$ and $\lambda_+(\varepsilon, E)$. This assumption is not restrictive as simplicity for these extremal eigenvalues is a generic property. Indeed we observe simplicity in all our numerical tests.

5.2 The Constrained Gradient System and its Stationary Points

In this section we are deriving from the free gradient determined in Theorem 3 the constrained gradient of the considered functional, that is the projected gradient (with respect to the Frobenius inner product) onto the manifold $\Omega \cap \Pi_\varepsilon$, which consists of perturbations E of unit norm which preserve the structure of W .

In order to obtain the constrained gradient system, we need to project the unconstrained gradient given by Theorem 3 onto the feasible set and also to normalize E to preserve its unit norm. Using the Karush-Kuhn-Tucker conditions on a time interval where the set of 0-weight edges remain unchanged, the projection is done via the mapping $\mathbb{P}_+ G(\varepsilon, E)$, where

$$[\mathbb{P}_+ X]_{ij} = \begin{cases} X_{ij}, & [W_1 + \varepsilon E]_{ij} > 0 \\ 0, & \text{otherwise} \end{cases}.$$

Further, in order to comply with the constraint $\|E(t)\|^2 = 1$, we must have

$$0 = \frac{1}{2} \frac{d}{dt} \|E(t)\|^2 = \langle E(t), \dot{E}(t) \rangle. \quad (16)$$

Thus, we obtain the following constrained optimization problem for the admissible direction of the steepest descent

Lemma 3 (Direction of steepest admissible descent) Let $E, G \in \mathbb{R}^{m \times m}$ with G given by (13), and $\|E\| = 1$. On a time interval where the set of 0-weight edges remains unchanged, the gradient system reads

$$\dot{E}(t) = -\mathbb{P}_+ G(\varepsilon, E(t)) + \kappa \mathbb{P}_+ E(t), \quad \text{where } \kappa = \frac{\langle \varepsilon, G(E(t)), \mathbb{P}_+ E(t) \rangle}{\|\mathbb{P}_+ E(t)\|^2}. \quad (17)$$

Proof We need to orthogonalize $\dot{E}(t)$ with respect to $E(t)$. To this end, we introduce a linear orthogonality correction, i.e. we set $\dot{E} = \mathbb{P}_+ (-G - \kappa E)$, and we determine κ by imposing the constraint $\langle E, \dot{E} \rangle = 0$. We then have

$$0 = \langle E, \dot{E} \rangle = \langle E, \mathbb{P}_+ (-G - \kappa E) \rangle = -\langle \mathbb{P}_+ E, G \rangle - \kappa \langle \mathbb{P}_+ E, \mathbb{P}_+ E \rangle,$$

and the result follows. □

Equation (17) suggests that the system goes “primarily” in the direction of the antigradient $-G(\varepsilon, E)$, thus the functional is expected to decrease along it.

Lemma 4 (Monotonicity) *Let $E(t)$ of unit Frobenius norm satisfy the differential equation (17), with G given by (13). Then, $F(\varepsilon, E(t))$ decreases monotonically with t .*

Proof We consider first the simpler case where the non-negativity projection does not apply so that $G = G(\varepsilon, E)$ (without $\mathbb{P}_+$). Then

$$\begin{aligned} \frac{d}{dt} F(\varepsilon, E)(t) &= \left\langle \nabla_E F_k(\varepsilon, E), \dot{E} \right\rangle = \langle \varepsilon G(\varepsilon, E(t)), -G(\varepsilon, E(t)) + \kappa E(t) \rangle \\ &= -\varepsilon \|G(\varepsilon, E)\|^2 + \varepsilon \frac{\langle G(\varepsilon, E), E \rangle}{\langle E, E \rangle} \langle G(\varepsilon, E), E \rangle \\ &= \varepsilon \left(-\|G(\varepsilon, E)\|^2 + \frac{|\langle G(\varepsilon, E), E \rangle|^2}{\|E\|^2} \right) \leq 0 \end{aligned} \quad (18)$$

where the final estimate is given by the Cauchy-Bunyakovsky-Schwarz inequality. The derived inequality holds on the time interval without the change in the support of $\mathbb{P}_+$ (so that no new edges are prohibited by the non-negativity projection). $\square$

5.3 Free Gradient Transition in the Outer Level

The optimization problem in the *inner level* is non-convex due to the non-convexity of the functional $F(\varepsilon, E)$. Hence, for a given ε , the computed minimizer $E(\varepsilon)$ may depend on the initial guess $E_0 = E_0(\varepsilon)$.

The effects of the initial choice are particularly important upon the transition $\widehat{\varepsilon} \rightarrow \varepsilon = \widehat{\varepsilon} + \Delta\varepsilon$ between constrained inner levels: given reasonably small $\Delta\varepsilon$, one should expect relatively close optimizers $E(\widehat{\varepsilon})$ and $E(\varepsilon)$, and, hence, the initial guess $E_0(\varepsilon)$ being close to and dependent on $E(\varepsilon)$.

This choice, which seems very natural, determines however a discontinuity

$$F(\widehat{\varepsilon}, E(\widehat{\varepsilon})) \neq F(\varepsilon, E(\widehat{\varepsilon})),$$

which may prevent the expected monotonicity property with respect to ε in the (likely unusual case) where $F(\widehat{\varepsilon}, E(\widehat{\varepsilon})) < F(\varepsilon, E(\widehat{\varepsilon}))$. This may happen in particular when $\Delta\varepsilon$ is not taken small; since the choice of $\Delta\varepsilon$ is driven by a Newton-like iteration we are interested in finding a way to prevent this situation and making the whole iterative method more robust. The goal is that of guaranteeing monotonicity of the functional both with respect to time and with respect to ε (Fig. 5).

When in the outer iteration we increase ε from a previous value $\widehat{\varepsilon} < \varepsilon$, we have the problem of choosing a suitable initial value for the constrained gradient system (17), such that at the stationary point $E(\widehat{\varepsilon})$ we have $F(\widehat{\varepsilon}, E(\widehat{\varepsilon})) < F(\varepsilon, E(\widehat{\varepsilon}))$ (which we assume both positive, that is on the left of the value ε^* which identifies the closest zero of the functional).

In order to maintain monotonicity with respect to time and also with respect to ε , it is convenient to start to look at the optimization problem with value ε , with the initial datum $\delta W_1 = \widehat{\varepsilon} E(\widehat{\varepsilon})$ of norm $\widehat{\varepsilon} < \varepsilon$.

This means we have essentially to optimize with respect to the inequality constraint $\|\delta W_1\| \leq \varepsilon$, or equivalently solve the problem (now with inequality constrain on $\|E\|_F$):

$$E(\varepsilon) = \arg \min_{E \in \Omega, \|E\| \leq 1} F(\varepsilon, E)$$

The situation changes only slightly from the one discussed above. If $\|E\| < 1$, every direction is admissible, and the direction of the steepest descent is given by the negative




Fig. 5 The scheme of alternating constrained (blue, $\|E(t)\| \equiv 1$) and free gradient (red) flows. Each stage inherits the final iteration of the previous stage as initial $E_0(\varepsilon_i)$ or $\tilde{E}_0(\varepsilon_i)$ respectively; constrained gradient is integrated till the stationary point ($\|\nabla F(E)\| = 0$), free gradient is integrated until $\|\delta W_1\| = \varepsilon_i + \Delta\varepsilon$. The scheme alternates until the target functional vanishes ($F(\varepsilon, E) = 0$) (Color figure online)

gradient. So we choose the free gradient flow

$$\dot{E} = -\mathbb{P}_+ G(\varepsilon, E(t)) \quad \text{as long as } \|E(t)\| < 1. \quad (19)$$

When $\|E(t)\| = 1$, then there are two possible cases. If $\langle \mathbb{P}_+ G(\varepsilon, E), E \rangle \geq 0$, then the solution of (19) has

$$\frac{d}{dt} \|E(t)\|^2 = 2 \langle \dot{E}, E \rangle = -2 \langle \mathbb{P}_+ G(\varepsilon, E(t)), E \rangle \leq 0,$$

and hence the solution of (19) remains of Frobenius norm at most 1.

Otherwise, if $\langle \mathbb{P}_+ G(\varepsilon, E), E \rangle < 0$, the admissible direction of steepest descent is given by the right-hand side of (17), and so we choose this ODE to evolve E . The situation can be summarized as taking, if $\|E(t)\| = 1$,

$$\dot{E} = -\mathbb{P}_+ G(\varepsilon, E) + \mu E \quad \text{with } \mu = \min(0, \kappa) \quad (20)$$

with $\kappa = \langle G(\varepsilon, E), \mathbb{P}_+ E \rangle / \|\mathbb{P}_+ E\|^2$. Along the solutions of (20), the functional F decays monotonically, and stationary points of (20) (i.e. points such that $\dot{E} = 0$) with $\mathbb{P}_+ G(\varepsilon, E(t)) \neq 0$ are characterized by

$$E \text{ is a negative real multiple of } \mathbb{P}_+ G(\varepsilon, E(t)). \quad (21)$$

If it can be excluded that the gradient $\mathbb{P}_+ G(\varepsilon, E(t))$ vanishes at an optimizer, it can thus be concluded that the optimizer of the problem with inequality constraints is a stationary point of the gradient flow (17) for the problem with equality constraints.

Remark 3 As a result, $F(\varepsilon, E(t))$ monotonically decreases both with respect to time t and to the value of the norm ε , when $\varepsilon \leq \varepsilon^*$.

6 Algorithm Details

In Algorithm 1 we provide the pseudo-code of the whole bi-level procedure. The initial “ α -phase” is used to choose an appropriate value for the regularization parameter α . In order

Algorithm 1 Pseudo-code of the complete constrained- and free-gradient flow.

Require: initial edge perturbation guess E_0 ; initial $\varepsilon_0 > 0$; ε -stepsize $\Delta\varepsilon > 0$; bounds α_* , α^* for the α -phase;

```

1:  $\alpha, E \leftarrow \text{AlphaPhase}(E_0, \varepsilon_0, \alpha_*, \alpha^*)$  ▷ for details see Section 7
2: while  $|F(\varepsilon, E)| < 10^{-6}$  do
3:    $\varepsilon \leftarrow \varepsilon + \Delta\varepsilon$ 
4:    $E \leftarrow \frac{\varepsilon}{\varepsilon + \Delta\varepsilon} E$  ▷ before the free gradient  $\|E\| < 1$ 
5:    $E_i \leftarrow \text{FreeGradientFlow}(E, \Delta\varepsilon, \varepsilon)$  ▷ see Section 5.3
6:    $E \leftarrow \text{ConstrainedGradientFlow}(E, \varepsilon)$  ▷ see Section 6
7: end while

```

to avoid the case in which the penalizing term on the right-hand side of (10) dominates the loss $F(\varepsilon, E(t))$ in the early stages of the descent flow, we select α by first running such an initial phase, prior to the main alternated constrained/free gradient loop. In this phase, we fix a small $\varepsilon = \varepsilon_0$ and run the constrained gradient integration for an initial $\alpha = \alpha_*$. After the computation of a local optimum E_* , we then increase α and rerun for the same ε_0 with E_* as starting point. We iterate until no change in E_* is observed or until α reaches an upper bound α^* .

The resulting value of α and E_* are then used in the main loop where we first increase ε by the chosen step size, we rescale E_i by $0 < \varepsilon/(\varepsilon + \Delta\varepsilon) < 1$, and then we perform the free gradient integration described in Sect. 5.3 until we reach a new point E_i on the unit sphere $\|E_i\| = 1$. Then, we perform the inner constrained gradient step by integrating Equation (17), iterating the following two-step norm-corrected Euler scheme:

$$\begin{cases} E_{i+1/2} = E_i - h_i (\mathbb{P}_+ G(E_i, \varepsilon) - \kappa_i \mathbb{P}_+ E_i) . \\ E_{i+1} = \mathbb{P}_{\Pi_\varepsilon} E_{i+1/2} / \|\mathbb{P}_{\Pi_\varepsilon} E_{i+1/2}\| \end{cases} \quad (22)$$

where the second step is necessary to numerically guarantee the Euler integration remains in the set of admissible flows since the discretization does not conserve the norm and larger steps h_i may violate the non-negativity of the weights.

In both the free and constrained integration phases, since we aim to obtain the solution at $t \rightarrow \infty$ instead of the exact trajectory, we favor larger steps h_i given that the established monotonicity is conserved. Specifically, if $F(\varepsilon, E_{i+1}) < F(\varepsilon, E_i)$, then the step is *accepted* and we set $h_{i+1} = \beta h_i$ with $\beta > 1$; otherwise, the step is *rejected* and repeated with a smaller step $h_i \leftarrow h_i/\beta$.

Remark 4 The step acceleration strategy described above, where βh_i is immediately increased after one accepted step, may lead to “oscillations” between accepted and rejected steps in the event the method would prefer to maintain the current step size h_i . To avoid this potential issue, in our experiments we actually increase the step length after two consecutive accepted steps. Alternative step-length selection strategies are also possible, for example, based on Armijo’s rule or non-monotone stabilization techniques [18].

6.1 Computational Costs

Each step of either the free or the constrained flows requires one step of explicit Euler integration along the anti-gradient $-\nabla_E F(\varepsilon, E(t))$. As discussed in Sect. 5, the construction of such a gradient requires several sparse and diagonal matrix–vector multiplications as well as the computation of the smallest nonzero eigenvalue of both $\bar{L}_1^{up}(\varepsilon, E)$ and $\bar{L}_0(\varepsilon, E)$. The latter two represent the major computational requirements of the numerical procedure.

Fortunately, as both matrices are of the form $A^\top A$, with A being either of the two boundary or co-boundary operators $\bar{B}_2$ and $\bar{B}_1^\top$, we have that both the two eigenvalue problems boil down to a problem of the form

$$\min_{x \perp \ker A} \frac{\|Ax\|}{\|x\|}$$

i.e., the computation of the smallest singular value of the sparse matrix A . This problem can be approached by a sparse singular value solver based on a Krylov subspace scheme for the pseudo inverse of $A^\top A$. In practice, we implement the pseudo inversion by solving the corresponding least squares problems

$$\min_x \|\bar{L}_1^{up}(\varepsilon, E)x - b\|, \quad \min_x \|\bar{L}_0(\varepsilon, E)x - b\|,$$

which, in our experiments, we solved using the least square minimal-residual method (LSMR) from [15]. This approach allows us to use a preconditioner for the normal equation corresponding to the least square problem. For simplicity, in our tests we used a constant preconditioner computed by means of an incomplete Cholesky factorization of the initial unperturbed $\bar{L}_1^{up}$, or $\bar{L}_0$. Possibly, much better performance can be achieved with a tailored preconditioner that is efficiently updated throughout the matrix flow. We also note that the eigenvalue problem for the graph Laplacian $\bar{L}_0(\varepsilon, E)$ may be alternatively approached by a combinatorial multigrid strategy [30]. However, designing a suitable preconditioning strategy goes beyond the scope of this work and will be the subject of future investigations.

7 Numerical Experiments

In this section, we provide several synthetic and real-world example applications of the proposed stability algorithm. The code for all the examples is available at <https://github.com/COMPILELab/HOLaGraF>. All experiments are run until the global stopping criterion $|F(\varepsilon, E(t))| < 10^{-6}$ is met. The parameters μ and α are chosen as follows. Concerning μ , we set $\mu = 0.75\mu_2$, where μ_2 is the smallest nonzero eigenvalue of the initial Laplacian $\bar{L}_0$. As for α , we run the α -phase described in Sect. 6 with parameters $\varepsilon_0 = 10^{-3}$, $\alpha_* = 1$ and $\alpha^* = 100$.

7.1 Illustrative Example

We consider here a small example of a simplicial complex $\mathcal{K}$ of order 2 consisting of eight 0-simplices (vertices), twelve 1-simplices (edges), four 2-simplices $\mathcal{V}_2 = \{[1, 2, 3], [1, 2, 8], [4, 5, 6], [5, 6, 7]\}$ and one corresponding hole $[2, 3, 4, 5]$, hence, $\beta_1 = 1$. By design, the dimensionality of the homology group $\bar{\mathcal{H}}_1$ can be increased only by eliminating edges $[1, 2]$ or $[5, 6]$; for the chosen weight profile $w_1([1, 2]) > w_1([5, 6])$, hence, the method should converge to the minimal perturbation norm $\varepsilon = w_1([5, 6])$ by eliminating the edge $[5, 6]$, Fig. 6.

The exemplary run of the optimization framework in time is shown on Fig. 7. The top panel of Fig. 7 provides the continued flow of the target functional $F(\varepsilon, E(t))$ consisting of the initial α -phase (in green) and alternated constrained (in blue) and free gradient (in orange) stages. As stated above, $F(\varepsilon, E(t))$ is strictly monotonic along the flow since the support of $\mathbb{P}_+$ does not change. Since the initial setup is not pathological with respect to the connectivity, the initial α -phase essentially reduces to a single constrained gradient flow and




Fig. 6 Simplicial complex $\mathcal{K}$ on 8 vertices for the illustrative run (on the left): all 2-simplices from $\mathcal{V}_2$ are shown in blue, the weight of each edge $w_1(e_i)$ is given on the figure. On the right: perturbed simplicial complex $\mathcal{K}$ through the elimination of the edge $[5, 6]$ creating additional hole $[5, 6, 7, 8]$




Fig. 7 Illustrative run of the framework determining the topological stability: the top pane—the flow of the functional $F(\varepsilon, E(t))$; the second pane—the flow of $\sigma(\bar{L}_1)$, λ_+ is highlighted; third pane — the change of the perturbation norm $\|E(t)\|$; the bottom pane—the heatmap of the perturbation profile $E(t)$

terminates after one run with $\alpha = \alpha_*$. The constrained gradient stages are characterized by a slow changing $E(t)$, which is essentially due to the flow performing small adjustments to find the correct rotation on the unit sphere, whereas the free gradient stage quickly decreases the perturbation.

The second panel shows the behaviour of first non-zero eigenvalue $\lambda_+(\varepsilon, E(t))$ (solid line) of $\bar{L}_1^{up}(\varepsilon, E(t))$ dropping through the ranks of $\sigma(\bar{L}_1(\varepsilon, E(t)))$ (semi-transparent); similar to the case of the target functional $F(\varepsilon, E(t))$, $\lambda_+(\varepsilon, E(t))$ monotonically decreases. The rest of the eigenvalues exhibit only minor changes, and the rapidly changing λ_+ successfully passes through the connectivity threshold μ (dotted line).




Fig. 8 Benchmarking Results on the Synthetic Triangulation Dataset: varying sparsities $\nu = 0.35, 0.5$ and $N = 16, 22, 28, 34, 40$; each network is sampled 10 times. Shapes correspond to the number of eliminated edges in the final perturbation: 1 : $\circ$, 2 : $\square$, 3 : $\diamond$, 4 : $\triangle$. For each pair (ν, N) , the un-preconditioned and Cholesky-preconditioned execution times are shown

The third and the fourth panels show the evolution of the norm of the perturbation $\|E(t)\|$ and the perturbation $E(t)$ itself, respectively. The norm $\|E(t)\|$ is conserved during the constrained-gradient and the α - stages; these stages correspond to the optimization of the perturbation shape, as shown by the small positive values at the beginning of the bottom panel which eventually vanish. During the free gradient integration the norm $\|E(t)\|$ increases, but the relative change of the norm declines with the growth of ε_i to avoid jumping over the smallest possible ε . Finally, due to the simplicity of the complex, the edge we want to eliminate, 56, dominates the flow from the very beginning (see bottom panel); such a clear pattern persists only in small examples, whereas for large networks the perturbation profile is initially spread out among all the edges.

7.2 Triangulation Benchmark

To provide more insight into the computational behavior of the method, we synthesize here an almost planar graph dataset. Namely, we assume N uniformly sampled vertices on the unit square with a network built by the Delaunay triangulation; then, edges are randomly added or erased to obtain the sparsity ν (so that the graph has $\frac{1}{2}\nu N(N-1)$ edges overall). An order-2 simplicial complex $\mathcal{K} = (\mathcal{V}_0, \mathcal{V}_1, \mathcal{V}_2)$ is then formed by letting $\mathcal{V}_0$ be the generated vertices, $\mathcal{V}_1$ the edges, and $\mathcal{V}_2$ every 3-clique of the graph; edges' weights are sampled uniformly between $1/4$ and $3/4$, namely $w_1(e_i) \sim U[\frac{1}{4}, \frac{3}{4}]$.

An example of such triangulation is shown in Fig. 8a; here $N = 8$ and edges $[6, 8]$ and $[2, 7]$ were eliminated to achieve the desired sparsity.

We sample networks with a varying number of vertices $N = 10, 16, 22, 28$ and varying sparsity pattern $\nu = 0.35, 0.5$ which determine the number of edges in the output as $m = \nu \frac{N(N-1)}{2}$. Due to the highly randomized procedure, topological structures of a sampled graph with a fixed pair of parameters may differ substantially, so 10 networks with the same (N, ν) pair are generated. For each network, the working time (without considering the sampling itself) and the resulted perturbation norm ε , and are reported in Fig. 8b and Fig. 8c, respectively. As anticipated in Sect. 6.1, we show the performance of two implementations of the method, one based on LSMR and one based on LSMR preconditioned by using the incomplete Cholesky factorization of the initial matrices. We observe that,

- the computational cost of the whole procedure lies between $\mathcal{O}(m^2)$ and $\mathcal{O}(m^3)$
- denser structures, with a higher number of vertices, result in the higher number of edges being eliminated; at the same time, even most dense cases still can exhibit structures requiring the elimination of a single edge, showing that the flow does not necessarily favor multi-edge optima;
- the required perturbation norm ε is growing with the size of the graph, Fig. 8c, but not too fast: it is expected that denser networks would require larger ε to create a new hole; at the same time if the perturbation were to grow drastically with the sparsity ν , it would imply that the method tries to eliminate sufficiently more edges, a behavior that resembles convergence to a sub-optimal perturbation;
- preconditioning with a constant incomplete Cholesky multiplier, computed for the initial Laplacians, provides a visible execution time gain for medium and large networks. Since the quality of the preconditioning deteriorates as the flow approaches the minimizer (as a non-zero eigenvalue becomes 0), it is worth investigating the design of a preconditioner for the up-Laplacian that can be efficiently updated.

7.3 Transportation Networks

Finally, we provide an application to real-world examples based on city transportation networks. We consider networks for Bologna, Anaheim, Berlin Mitte, and Berlin Tiergarten; each network consists of nodes—intersections/public transport stops—connected by edges (roads) and subdivided into zones; for each road the free flow time, length, speed limit are known; moreover, the travel demand for each pair of nodes is provided through the dataset of recorded trips. All the datasets used here are publicly available at <https://github.com/bstabler/TransportationNetworks>; Bologna network is provided by the Physic Department of the University of Bologna (enriched through the Google Maps API <https://developers.google.com/maps>).

The regularity of city maps naturally lacks 3-cliques, hence forming the simplicial complex based on triangulations as done before frequently leads to trivial outcomes. Instead, here we “lift” the network to city zones, thus more effectively grouping the nodes in the graph. Specifically:

1. We consider the completely connected graph where the nodes are zones in the city/region;
2. The free flow time between two zones is temporarily assigned as a weight of each edge: the time is as the shortest path between the zones (by the classic Dijkstra algorithm) on the initial graph;
3. Similarly to what is done in the filtration used for persistent homology, we filter out excessively distant nodes; additionally, we exclude the longest edges in each triangle in case it is equal to the sum of two other edges (so the triangle is degenerate and the trip by the longest edge is always performed through to others);
4. finally, we use the travel demand as an actual weight of the edges in the final network; travel demands are scaled *logarithmically* via the transformation $w_i \mapsto \log_{10} \left(\frac{w_i}{0.95 \min w_i} \right)$; see the example on the left panel of Fig. 9.

Given the definition of weights in the network, high instability (corresponding to small perturbation norm ε) implies structural phenomena around the “almost-hole”, where the faster and shorter route is sufficiently less demanded.

In the case of Bologna, Fig. 9, the algorithm eliminates the edge [11, 47] (Casalecchio di Reno – Pianoro) creating a new hole $6 - 11 - 57 - 47$. We also provide examples of




Fig. 9 Example of the Transportation Network for Bologna. Left pane: original zone graph where the width of edges corresponds to the weight, to-be-eliminated edge is colored in red. Right pane: eigenflows, original and created; color and width correspond to the magnitude of entries (Color figure online)

Table 1 Topological instability of the transportation networks: filtered zone networks with the corresponding perturbation norm ε and its percentile among $w_1(\cdot)$ profile

	Network			β_1	Logarithmic weights		
	n	m	Δ		Time	ε	p
Bologna	60	175	171	2	2.43 s	0.65	0.003
					[11, 47] (4th smallest)		
Anaheim	38	159	221	1	5.39 s	0.57	0.003
					[10, 29] (11th smallest)		
Berlin-Tiergarten	26	63	55	0	2.46 s	1.18	0.015
					[6, 16] (20th smallest)		
Berlin-Mitte	98	456	900	1	127 s	0.887	0.0016
					[57, 87] (6th), [58, 87], (17th)		

For each simplicial complex the number of nodes, edges and triangles in $\mathcal{V}_2(\mathcal{K})$ are provided alongside the initial number of holes β_1 . The results of the algorithm consist of the perturbation norm, ε , computation time, and approximate percentile p

the eigenflows in the kernel of the Hodge Laplacian (original and additional perturbed): original eigenvectors correspond to the circulations around holes 7 – 26 – 12 – 20 and 8 – 21 – 20 – 16 – 37 non-locally spread in the neighborhood [29].

The results for four different networks are summarized in the Table 1; p mimics the percentile, $\varepsilon / \sum_{e \in \mathcal{V}_1} w_i(e)$, showing the overall small perturbation norm contextually. At the same time, we emphasize that except Bologna (which is influenced by the geographical topology of the land), the algorithm does not choose the smallest weight possible; indeed, given our interpretation of the topological instability, the complex for Berlin-Tiergarten is stable and the transportation network is effectively constructed.

8 Discussion

In the current work, we formulate the notion of k -th order topological stability of a simplicial complex $\mathcal{K}$ as the smallest data perturbation required to create one additional k -th order hole in $\mathcal{K}$. By introducing an appropriate weighting and normalization, the stability is reduced to a matrix nearness problem solved by a bi-level optimization procedure. Despite the highly nonconvex landscape, our proposed procedure alternating constrained and free gradient steps yields a monotonic descending scheme. Our experiments show that this approach is generally successful in computing the minimal perturbation increasing $\beta_1(\varepsilon, E)$, even for potentially difficult cases, as the one proposed in Sect. 7.1.

For simplicity, here we limit our attention to the smallest perturbation that introduces only one new hole. However, a simple modification may be employed to address the case of the introduction of m new holes: include the sum of m nonzero eigenvalues of $L_1^{up}(\varepsilon, E)$ rather than just the first one in the spectral functional (10). We also remark that, due to the spectral inheritance principle 1, the proposed framework for $\mathcal{H}_1$ can be in principle extended to a general $\mathcal{H}_k$; however, this extension requires nontrivial considerations on the data modification procedure and on the numerical linear algebra tools, as a nontrivial topology of higher-order requires a much denser network.

Different improvements are possible in terms of numerical implementation, including investigating the use of more sophisticated (e.g. implicit) integrators for the gradient flow system (17), which would additionally require the use of higher-order derivatives of $\lambda_+(\varepsilon, E)$. Moreover, as already mentioned in Sect. 6.1, the numerical method for the computation of the small singular values would benefit from the use of an efficient preconditioner that can be effectively updated throughout the flow. Investigations in this direction are in progress and will be the subject of future work.

Funding Open access funding provided by Gran Sasso Science Institute - GSSI within the CRUI-CARE Agreement. Nicola Guglielmi acknowledges that his research was supported by funds from the Italian MUR (Ministero dell'Università e della Ricerca) within the PRIN 2017 Project “Discontinuous dynamical systems: theory, numerics and applications”. Authors are affiliated to the INdAM-GNCS (Gruppo Nazionale di Calcolo Scientifico) and acknowledge support from MUR’s Dipartimento di Eccellenza program grant on “Pattern Analysis and Engineering”.

Data availability Enquiries about data availability should be directed to the authors.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Altenburger, K.M., Ugander, J.: Monophily in social networks introduces similarity among friends-of-friends. *Nat. Hum. Behav.* **2**(4), 284–290 (2018)
2. Andreotti, E., Edelmann, D., Guglielmi, N., Lubich, C.: Constrained graph partitioning via matrix differential equations. *SIAM J. Matrix Anal. Appl.* **40**(1), 1–22 (2019). <https://doi.org/10.1137/17M1160987>
3. Andreotti, E., Edelmann, D., Guglielmi, N., Lubich, C.: Measuring the stability of spectral clustering. *Linear Algebra Appl.* **610**, 673–697 (2021). <https://doi.org/10.1016/j.laa.2020.10.015>
4. Arrigo, F., Higham, D.J., Tudisco, F.: A framework for second-order eigenvector centralities and clustering coefficients. *Proc. R. Soc. A* **476**(2236), 20190724 (2020)
5. Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.-G., Petri, G.: Networks beyond pairwise interactions: structure and dynamics. *Phys. Rep.* **874**, 1–92 (2020). <https://doi.org/10.1016/j.physrep.2020.05.004>
6. Benson, A.R.: Three hypergraph eigenvector centralities. *SIAM J. Math. Data Sci.* **1**(2), 293–312 (2019)
7. Benson, A.R., Gleich, D.F., Leskovec, J.: Higher-order organization of complex networks. *Science* **353**(6295), 163–166 (2016). <https://doi.org/10.1126/science.aad9029>
8. Benson, A.R., Abebe, R., Schaub, M.T., Jadbabaie, A., Kleinberg, J.: Simplicial closure and higher-order link prediction. *Proc. Natl. Acad. Sci.* **115**(48), E11221–E11230 (2018)
9. Bick, C., Gross, E., Harrington, H.A., Schaub, M.T.: What are higher-order networks? *SIAM Review.* (2023). [arXiv:2104.11329](https://arxiv.org/abs/2104.11329)
10. Chazal, F., De Silva, V., Oudot, S.: Persistence stability for geometric complexes. *Geom. Dedicata.* **173**(1), 193–214 (2014)
11. Chen, Y.-C., Meila, M.: The decomposition of the higher-order homology embedding constructed from the k -Laplacian. *Adv. Neural. Inf. Process. Syst.* **34**, 15695–15709 (2021)
12. Chen, Y.-C., Meilă, M., Kevrekidis, I.G.: Helmholtzian eigenmap: topological feature discovery and edge flow learning from point cloud data (2021). [arXiv:2103.07626](https://arxiv.org/abs/2103.07626). <https://doi.org/10.48550/ARXIV.2103.07626>
13. Ebli, S., Spreemann, G.: A notion of harmonic clustering in simplicial complexes. In: 2019 18th IEEE International Conference on Machine Learning and Applications (ICMLA), pp. 1083–1090. IEEE (2019)
14. Fiedler, M.: Laplacian of graphs and algebraic connectivity. *Banach Cent. Publ.* **25**(1), 57–70 (1989)
15. Fong, D.C.-L., Saunders, M.: LSMR: an iterative algorithm for sparse least-squares problems. *SIAM J. Sci. Comput.* **33**(5), 2950–2971 (2011)
16. Fountoulakis, K., Li, P., Yang, S.: Local hyper-flow diffusion. *Adv. Neural. Inf. Process. Syst.* **34**, 27683–27694 (2021)
17. Gautier, A., Tudisco, F., Hein, M.: Nonlinear Perron–Frobenius theorems for nonnegative tensors. *SIAM Rev.* **65**(2), 495–536 (2023)
18. Grippo, L., Lampariello, F., Lucidi, S.: A class of nonmonotone stabilization methods in unconstrained optimization. *Numer. Math.* **59**(1), 779–805 (1991)
19. Guglielmi, N., Lubich, C.: Matrix nearness problems and eigenvalue optimization. in preparation (2022)
20. Horak, D., Jost, J.: Spectra of combinatorial Laplace operators on simplicial complexes. *Adv. Math.* **244**, 303–336 (2013)
21. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press, Cambridge (1990)
22. Lim, L.-H.: Hodge Laplacians on graphs. *SIAM Rev.* **62**(3), 685–715 (2015)
23. Muhammad, A., Egerstedt, M.: Control using higher order Laplacians in network topologies. In: Proceedings of 17th international symposium on mathematical theory of networks and systems, pp. 1024–1038. Citeseer (2006)
24. Nettasinghe, B., Krishnamurthy, V., Lerman, K.: Diffusion in social networks: Effects of monophilic contagion, friendship paradox and reactive networks. *IEEE Trans. Netw. Sci. Eng.* **7**, 1121–1132 (2019)
25. Neuhäuser, L., Lambiotte, R., Schaub, M.T.: Consensus dynamics and opinion formation on hypergraphs. In: *Higher-Order Systems*, pp. 347–376. Springer (2022)
26. Neuhäuser, L., Scholkemper, M., Tudisco, F., Schaub, M.T.: Learning the effective order of a hypergraph dynamical system (2023). [arXiv preprint arXiv:2306.01813](https://arxiv.org/abs/2306.01813)
27. Otter, N., Porter, M.A., Tillmann, U., Grindrod, P., Harrington, H.A.: A roadmap for the computation of persistent homology. *EPJ Data Sci.* **6**, 1–38 (2017)
28. Qi, L., Luo, Z.: *Tensor Analysis*. Society for Industrial and Applied Mathematics, Philadelphia (2017)
29. Schaub, M.T., Benson, A.R., Horn, P., Lippner, G., Jadbabaie, A.: Random walks on simplicial complexes and the normalized Hodge 1-Laplacian. *SIAM Rev.* **62**(2), 353–391 (2020)
30. Spielman, D.A., Teng, S.-H.: Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM J. Matrix Anal. Appl.* **35**(3), 835–885 (2014)

31. Tudisco, F., Hein, M.: A nodal domain theorem and a higher-order cheeger inequality for the graph p -Laplacian (2016). [arXiv:1602.05567](https://arxiv.org/abs/1602.05567)
32. Tudisco, F., Higham, D.J.: Core-periphery detection in hypergraphs. *SIAM J. Math. Data Sci*
33. Tudisco, F., Arrigo, F., Gautier, A.: Node and layer eigenvector centralities for multiplex networks. *SIAM J. Appl. Math.* **78**(2), 853–876 (2018)
34. Tudisco, F., Benson, A.R., Prokophchik, K.: Nonlinear higher-order label spreading. In: *Proceedings of the Web Conference 2021*, pp. 2402–2413 (2021)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.