

Edge Contractions and Simplicial Homology

Tamal K. Dey* Anil N. Hirani† Bala Krishnamoorthy‡ Gavin Smith§

July 18, 2013

Abstract

We study the effect of edge contractions on simplicial homology because these contractions have turned out to be useful in various applications involving topology. It was observed previously that contracting edges that satisfy the so called *link condition* preserves homeomorphism in low dimensional complexes, and homotopy in general. But, checking the link condition involves computation in all dimensions, and hence can be costly, especially in high dimensional complexes. We define a weaker and more local condition called the *p-link* condition for each dimension p , and study its effect on edge contractions. We prove the following: (i) For homology groups, edges satisfying the p - and $(p - 1)$ -link conditions can be contracted without disturbing the p -dimensional homology group. (ii) For relative homology groups, the $(p - 1)$ -, and the $(p - 2)$ -link conditions suffice to guarantee that the contraction does not introduce any new class in any of the resulting relative homology groups, though some of the existing classes can be destroyed. Unfortunately, the surjection in relative homology groups does not guarantee that no new relative torsion is created. (iii) For torsions, edges satisfying the p -link condition alone can be contracted without creating any *new* relative torsion, and the p -link condition cannot be avoided in some cases. The results on relative homology and relative torsion are motivated by recent results on computing optimal homologous chains, which state that such problems can be solved by linear programming if the complex has no relative torsion. Edge contractions that do not introduce new relative torsions can safely be availed in these contexts.

1 Introduction

The study of edge contractions in the context of graph theory [23], especially in graph minor theory [16] have resulted into many noteworthy results. The extension of edge contractions to simplicial complexes where the structure not only has vertices and edges, but also higher dimensional simplices has also turned out to be beneficial for shape representation in graphics and visualization [14, 11] and recently in topological data analysis [2]. In this paper, we present several results about edge contractions in simplicial complexes that can benefit extraction of topological attributes from a shape or data representation.

Topological attributes such as the rank of the homology groups, also known as betti numbers, and cycles representing the homology classes carry important information about a shape. Naturally, efforts have ensued to compute them efficiently in various applications. Examples include computing topological features for low dimensional complexes in graphics, visualizations, and sensor networks [3, 7, 21, 24], and for higher dimensional ones in data analysis [12]. As the input sizes in these applications grow with the advances in

*Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA. Email: tamaldehy@cse.ohio-state.edu

†Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA. Email: hirani@illinois.edu

‡Department of Mathematics, Washington State University, Pullman, WA, USA. Email: bkrishna@math.wsu.edu

§Department of Mathematics, Washington State University, Pullman, WA, USA. Email: gsmith@math.wsu.edu

data generation technology, methods to speed up the homology computations become more demanding. For example, the Vietoris-Rips complex has been recognized as a versatile data structure for inferring homological attributes from point cloud data [12]. However, because of its inclusive nature, this complex tends to contain a large number of simplices and the computation becomes prohibitive when the input point set is large. One way to tackle this issue is to use *edge contractions* [8, 11] that contract edges and collapse other simplices as a result. Quite naturally, edge contractions have already been proposed to tame the size of large Rips complexes [2]. In a recent work on computing optimal homologous cycles (OHCP), we have shown how relative homology and torsions play a role in guaranteeing a polynomial time optimization [9]. To avail the benefit of edge contractions in this context we need to understand its effect on relative homology and torsion. Motivated by these applications, we make a systematic study of the effects of edge contractions on simplicial homology.

The effects of edge contractions on topology was initially studied by Walkup [22] for 3-manifolds and then by Dey et al. [8] for more general domains. They showed that an edge e in a 2-complex or a 3-manifold can be contracted while preserving a homeomorphism between the complexes before and after the edge contraction if e satisfies a local condition called the *link condition*. The condition, due to its locality, is easily checkable at least for low dimensional complexes. Attali, Lieutier, and Salinas [1] showed that the result can be extended to the entire class of finite simplicial complexes if only homotopy instead of homeomorphism needs to be preserved. Since homotopy equivalent spaces have isomorphic homology groups, link conditions also suffice to preserve homology groups. However, verification of the link condition requires checking it at every dimension, which becomes costly for higher dimensional simplicial complexes.

In this work, we extend the above results in two directions: (i) we study edge contractions not only for homology groups, but also for relative homology groups and torsion groups, and (ii) we define a weaker and even more local condition called the *p -link condition* for each dimension p and analyze its effect on homology, relative homology, and torsion. Specifically, we prove that an edge contraction cannot destroy any homology class in dimension p if the edge satisfies the p -link condition alone. Furthermore, no new class is introduced if the edge satisfies the $(p - 1)$ -link condition. For relative homology, we show that no new homology class in dimension p relative to a particular subcomplex is created if the edge satisfies the $(p - 1)$ - and the $(p - 2)$ -link conditions in that *subcomplex*. This result also implies that no new relative homology classes are generated in the contracted complex as long as the contracting edge satisfies the $(p - 1)$ - and $(p - 2)$ -link conditions in the *original complex*. Unfortunately, this surjectivity in relative homology classes does not mean surjectivity in *relative torsions*, that is, the torsion subgroups of relative homology groups. Of course, if the edge additionally satisfies the p -link condition in the subcomplex, we have isomorphisms in relative homology groups which necessarily mean that no new relative torsion is created. We strengthen the condition to require the p -link condition alone if only relative torsion is of interest, which is the case for OHCP. An example shows that one cannot avoid the p -link condition if one is to guarantee that no new relative torsion is introduced.

Our result on homology preservation under p - and $(p - 1)$ -link conditions can be used to compute the betti numbers of large complexes after edge contractions. They can also be availed to compute actual representative cycles in a small contracted complex, which can then be pulled back to the original complex through a systematic reversal of the contractions. Computations are saved by checking the link conditions for only a few dimensions instead of all dimensions. Similarly, our result for relative homology can be used to compute the betti numbers of the quotient complexes formed by the original complex relative to a subcomplex. Our result on torsion can be availed for computing the shortest cycle in a given homology class. This problem, termed the OHCP, is known to be NP-hard in general [5]. It has been recently shown that the OHCP is solvable in polynomial time when the homology is defined over \mathbb{Z} and the simplicial complex does not have relative torsion [9]. Similar results hold for related problems such as the optimal bounding chain problem (OBCP) [10] and the multiscale simplicial flat norm (MSFN) problem [15]. Therefore, edge contractions that preserve the absence of, or even better, eliminate relative torsion should be preferred. Fig-

ure 3 illustrates a case where an edge contraction does eliminate relative torsion and thus allows computing an optimal cycle by linear programming as shown in [9]. We use methods from algebraic topology to prove most of our results. Interestingly enough, the result on relative torsion requiring the p -link condition alone is proved only with graph theoretic techniques.

2 Background

We recall some basic concepts and definitions from algebraic topology relevant to our presentation. Refer to standard books, e.g., ones by Munkres [17] for details.

Definition 2.1. Given a vertex set V , a *simplicial complex* $K = K(V)$ is a collection of subsets $\{\sigma \subseteq V\}$ where $\sigma' \subseteq \sigma$ is in K if $\sigma \in K$. A subset $\sigma \in K$ of cardinality $p+1$ is called a p -*simplex*. If $\sigma' \subseteq \sigma$ ($\sigma' \subset \sigma$), we call σ' a *face* (*proper face*) of σ , and σ a *coface* (*proper coface*) of σ' . A map $h : K \rightarrow K'$ between two complexes is called *simplicial* if for every simplex $\{v_1, v_2, \dots, v_k\}$ in K , $\{h(v_1), h(v_2), \dots, h(v_k)\}$ is a simplex in K' .

An oriented simplex $\sigma = \{v_0, v_1, \dots, v_p\}$, also written as $v_0v_1 \dots v_p$, is an ordered set of vertices. The simplices σ_i with coefficients α_i in \mathbb{Z} can be added formally creating a chain $c = \sum_i \alpha_i \sigma_i$. These chains form the chain group C_p . The boundary $\partial_p \sigma$ of a p -simplex σ , $p \geq 0$, is the $(p-1)$ -chain that adds all the $(p-1)$ -faces of σ with orientation taken into consideration. This defines a boundary homomorphism $\partial_p : C_p \rightarrow C_{p-1}$. The kernel of ∂_p forms the p -cycle group $Z_p(K)$ and its image forms the $(p-1)$ -boundary group $B_{p-1}(K)$. The homology group $H_p(K)$ is the quotient group $Z_p(K)/B_p(K)$. Intuitively, a p -cycle is a collection of oriented p -simplices whose boundary is zero. It is a nontrivial cycle in H_p , if it is not a boundary of a $(p+1)$ -chain.

For a finite simplicial complex K , the groups of chains $C_p(K)$, cycles $Z_p(K)$, and $H_p(K)$ are all finitely generated abelian groups. By the fundamental theorem of finitely generated abelian groups [17, page 24] any such group G can be written as a direct sum of two groups $G = F \oplus T$ where $F \cong (\mathbb{Z} \oplus \dots \oplus \mathbb{Z})$ and $T \cong (\mathbb{Z}/t_1 \oplus \dots \oplus \mathbb{Z}/t_k)$ with $t_i > 1$ and t_i dividing t_{i+1} . The subgroup T is called the *torsion* of G . If $T = 0$, we say G is *torsion-free*.

A simplicial map $h : K \rightarrow K'$ between two simplicial complexes induces a homomorphism between their homology groups which we write as $h_* : H_p(K) \rightarrow H_p(K')$. The edge contractions that we deal with define such homomorphisms whose properties in relation to various p -link conditions are the focus of our work.

2.1 Link conditions

We first define edge contractions formally.

Definition 2.2. Let $ab = \{a, b\}$ be an edge in a simplicial complex K . An edge contraction of K is a surjective simplicial map $\gamma_{ab} : K \rightarrow K'$ induced by the vertex map $h : V(K) \rightarrow V(K')$ where h is identity everywhere except at b for which $h(b) = a$.

Authors of [8] investigated when such an edge contraction results in any change in the topology of the simplicial complex. They provided a sufficient condition termed the link condition, which guarantees that topology is preserved for certain simplicial complexes. This condition has been studied further by Attali, Lieutier, and Salinas [1] recently.

Definition 2.3. The *star* of a set $X \subseteq K$, denoted $\text{St } X$, is the set of cofaces of all $\sigma \in X$. For a subset S of K , the *closure* of S , denoted \overline{S} , is the set of simplices in S and all of their faces. Then the *link* of X , denoted $\text{Lk } X$, is the set of simplices in $\overline{\text{St } X}$ that do not belong to $\text{St } X$. In the left complex of Figure 1, the star of the edge ab consists of $ab, abd, abe, abde$. Its link is d, e, de .

The following definition of the link condition is taken from the work in [8], and we introduce a weaker condition called the p -link condition.

Definition 2.4. An edge $ab \in K$ satisfies the *link condition* in K if and only if $\text{Lk } a \cap \text{Lk } b = \text{Lk } ab$. It satisfies the p -link condition in K , if and only if either (i) $p \leq 0$, or (ii) $p > 0$ and every $(p - 1)$ -simplex $\xi \in \text{Lk } a \cap \text{Lk } b$ is also in $\text{Lk } ab$.

The p -link conditions are weaker than the link condition in the following sense.

Proposition 2.5. For any edge ab in a simplicial complex K , ab satisfies the link condition if and only if it satisfies the p -link condition for all $p \leq \dim(K)$.

Proof. We prove both directions by contrapositive. Assume ab does not satisfy the link condition. By the definition of star, $\text{St } ab = \text{St } a \cap \text{St } b$. Therefore, $\overline{\text{St } ab} \subset \overline{\text{St } a} \cap \overline{\text{St } b}$, and $\text{Lk } ab \subset \text{Lk } a \cap \text{Lk } b$. Therefore, there must be a simplex $\xi \in \text{Lk } a \cap \text{Lk } b$ where $\xi \notin \text{Lk } ab$. Let $p = \dim(\xi) + 1$. Then, ab does not satisfy the p -link condition. By the definition of link, there must be a p -simplex $\tau_1 \in K$ where $\tau_1 = \xi \cup a$, and also a p -simplex $\tau_2 \in K$ where $\tau_2 = \xi \cup b$. Therefore, $p \leq \dim(K)$.

Now assume ab does not satisfy the p -link condition for some $p \leq \dim(K)$. Then there is some $(p - 1)$ -simplex $\xi \in \text{Lk } a \cap \text{Lk } b$ that is not in $\text{Lk } ab$. Therefore, ab does not satisfy the link condition. \square

As an example, consider the two complexes in Figure 1. Both complexes have the lower tetrahedron adjoining the edge ab . The complex on the left has the top triangle missing. The link of the edge ab contains d, e, de in the left complex, and it satisfies the 2-link condition but not the 1-link condition. ab satisfies both the 1-link and the 2-link conditions in the complex on the right, since the link of ab contains c, d, e , and de .




Figure 1: p -link conditions for $p = 1, 2$.

3 Homology Preservation

In this section we study how the p -link conditions for edges being contracted affect the homology groups. Since the homology in dimension p depends only on the $(p - 1)$ -, p -, and $(p + 1)$ -skeleta, it is not surprising that only a subset of the relevant link conditions are sufficient to guarantee homology preservation. We make this intuition precise in the next theorem. Specifically, we prove that the p -link condition alone implies that no homology class in $H_p(K)$ is destroyed (injectivity) and the $(p - 1)$ -link condition alone implies that no new homology class is created (surjectivity).

Theorem 3.1. Let ab be an edge in a simplicial complex K and $\gamma_{ab} : K \rightarrow K'$ be an edge contraction. Then the induced homomorphism γ_{ab*} at the homology level has the following properties:

1. $\gamma_{ab*} : H_p(K) \rightarrow H_p(K')$ is injective if ab satisfies the p -link condition.
2. γ_{ab*} is surjective if ab satisfies the $(p - 1)$ -link condition.

To prove the above theorem, we use an intermediate complex \hat{K} constructed as follows. Let the cone $v * \sigma$ for a vertex $v \in K$ and a simplex $\sigma \in K$ be defined as the closure of the simplex $\sigma \cup \{v\}$. The cone $v * T$ for a subcomplex $T \subseteq K$ is the complex $v * T = \{v * \sigma \mid \sigma \in T\}$. We construct $\hat{K} = K \cup (a * \overline{\text{St } b})$.

In words, \hat{K} is constructed out of K by adding simplices that are obtained by coning from a to the closed star of b . First, observe that the $a * \overline{\text{St } b}$ in \hat{K} can deformation retract to $a * \text{Lk } b$ taking b to a , and we get K' . Therefore, we have the sequence $K \xrightarrow{i} \hat{K} \xrightarrow{r} K'$ where i and r are an inclusion and a deformation retraction, respectively, and $\gamma_{ab} = r \circ i$; see Figure 2 for an illustration. At the homology level we have the sequence of two homomorphisms where the one on right is an isomorphism induced by a deformation retract.

$$H_p(K) \xrightarrow{i_*} H_p(\hat{K}) \xrightarrow{r_*} H_p(K').$$

Since $\gamma_{ab*} = r_* \circ i_*$ at the homology level and r_* is an isomorphism, we have that γ_{ab*} is injective or surjective if and only if i_* is.




Figure 2: Edge contraction as a composition of an inclusion and a retraction.

Proposition 3.2. $i_* : H_p(K) \rightarrow H_p(\hat{K})$ is injective if ab satisfies the p -link condition, and is surjective if ab satisfies the $(p - 1)$ -link condition.

Proof. Since p -dimensional homology is determined only by the skeleton up to dimension $p + 1$, assume that K is only $(p + 1)$ -dimensional. Let S denote the subcomplex $a * \overline{\text{St } b}$. The Mayer-Vietoris sequence

$$H_p(K \cap S) \xrightarrow{(j_*, k_*)} H_p(K) \oplus H_p(S) \xrightarrow{(i_* - \ell_*)} H_p(\hat{K}) \xrightarrow{\pi} H_{p-1}(K \cap S)$$

is exact where i, j, k, ℓ are respective inclusion maps and π is the connecting homomorphism [13]. We examine the maps j_*, k_*, ℓ_* and prove the required properties for i_* .

First consider that ab satisfies the p -link condition. We examine the the complex $K \cap S$. It is the union of $\overline{\text{St } b}$ and the added simplices that are cones of a to the simplices in $\text{Lk } a \cap \text{Lk } b$. None of these added simplices can be p -dimensional since otherwise the p -link condition is violated. Any closed star has trivial homology, and thus $H_p(\text{St } b) = 0$. Since no simplex of dimension p or more is added to create $K \cap S$ from $\overline{\text{St } b}$, we still have $H_p(K \cap S) = 0$. Therefore, $\text{im } j_*$ and $\text{im } k_*$ are trivial. Because of the exactness of the sequence, the map $i_* - \ell_*$ is injective. However, ℓ_* is a zero map since $S = a * \overline{\text{St } b}$ has $H_p(S) = 0$. It follows that i_* is injective as we are required to prove.

Next, consider that ab satisfies the $(p - 1)$ -link condition. By the same logic as above, there is no $(p - 1)$ -simplex that can be added to $\overline{\text{St } b}$ to create $K \cap S$. Therefore, $H_{p-1}(K \cap S) = 0$ which implies that $\text{im } \pi$ is trivial. Because of the exactness, we have that $(i_* - \ell_*)$ is surjective which implies i_* is surjective as ℓ_* is a zero map. \square

4 Relative Homology Preservation

In this section we study how the p -link conditions affect the relative homology groups. The motivation comes from a recent result on the optimal homologous cycle problem (OHCP) whose efficient solution depends on relative torsion and hence on relative homology classes [9].

First, we set up some background. Let L_0 be a subcomplex of a simplicial complex L . The quotient group $C_p(L)/C_p(L_0)$ is called the group of *relative p -chains* of L modulo L_0 and is denoted $C_p(L, L_0)$. The boundary operator $\partial_p : C_p(L) \rightarrow C_{p-1}(L)$ and its restriction to L_0 induce a homomorphism

$$\partial_p^{(L, L_0)} : C_p(L, L_0) \rightarrow C_{p-1}(L, L_0).$$

Writing $Z_p(L, L_0) = \ker \partial_p^{(L, L_0)}$ for *relative cycles* and $B_{p-1}(L, L_0) = \text{im } \partial_p^{(L, L_0)}$ for *relative boundaries*, we obtain the *relative homology group* $H_p(L, L_0) = Z_p(L, L_0)/B_p(L, L_0)$.

A *pure simplicial complex* of dimension p is formed by a collection of p -simplices and their proper faces. We consider relative homology groups $H_p(L, L_0)$ where $L \subseteq K$ and $L_0 \subset L$ are pure subcomplexes of dimensions $(p + 1)$ and p , respectively.

To study the effect of edge contraction on relative homology groups, we note that the simplicial map $\gamma_{ab} : K \rightarrow K'$ naturally extends to pairs as $\gamma_{ab} : (L, L_0) \rightarrow (L', L'_0)$ where $\gamma_{ab}(L) = L'$ and $\gamma_{ab}(L_0) = L'_0$. Notice that since L is a subcomplex of K , it may not satisfy some link conditions even if K does. Therefore, it is possible a priori that relative homology groups may not be controlled by putting conditions on the edges being contracted in K . Nevertheless, we show that for every pair of subcomplexes (L', L'_0) in K' , there is a pair of subcomplexes (L, L_0) in K such that $\gamma_{ab}(L, L_0) = (L', L'_0)$ and $\gamma_{ab*} : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is surjective if ab satisfies the $(p - 1)$ -, and the $(p - 2)$ -link conditions in K . This implies that no new relative homology classes are created by such an edge contraction. In contrast, however, relative homology classes can be killed by edge contractions even if it satisfies all link conditions. We develop these results now.

The following result follows from the exact sequence of the pair [19, Theorem 5.8] and the *five lemma* in algebraic topology [19, Theorem 5.10]. It lets us connect our result for homology to relative homology groups.

Theorem 4.1 ([19]). *Let $h : (L, L_0) \rightarrow (L', L'_0)$ be a simplicial map. Let $f_i : H_i(L) \rightarrow H_p(L')$ and $g_i : H_i(L_0) \rightarrow H_p(L'_0)$ be the homomorphisms induced by h in the absolute homology groups for $i = p, p - 1$. The following diagram commutes where the horizontal sequences are exact.*

$$\begin{array}{ccccccccc} H_p(L_0) & \longrightarrow & H_p(L) & \longrightarrow & H_p(L, L_0) & \longrightarrow & H_{p-1}(L_0) & \longrightarrow & H_{p-1}(L) \\ \downarrow g_p & & \downarrow f_p & & \downarrow h & & \downarrow g_{p-1} & & \downarrow f_{p-1} \\ H_p(L'_0) & \longrightarrow & H_p(L') & \longrightarrow & H_p(L', L'_0) & \longrightarrow & H_{p-1}(L'_0) & \longrightarrow & H_{p-1}(L') \end{array}$$

Then, the following statements hold due to the five lemma:

- (i) *If f_p and g_{p-1} are surjective, and f_{p-1} is injective, then $h_* : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is surjective.*
- (ii) *If f_p and g_{p-1} are injective, and g_p is surjective, then $h_* : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is injective.*
- (iii) *If f_i and g_i are isomorphisms for $i = p, p - 1$, then $h_* : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is an isomorphism.*

If we take the map h to be the restriction of the simplicial map γ_{ab} to L and its subcomplex L_0 , we can use Theorem 3.1 to arrive at the following result:

Theorem 4.2. *Let $L \subseteq K$ be any pure subcomplex of dimension $p + 1$ and $L_0 \subset L$ be any of its pure subcomplexes of dimension p . Let $(L', L'_0) = \gamma_{ab}(L, L_0)$. Then, the following hold:*

- (i) *If ab satisfies the $(p - 2)$ - and $(p - 1)$ -link conditions in L_0 , then $\gamma_{ab*} : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is surjective;*

- (ii) If ab satisfies the $(p - 1)$ - and p -link conditions in L_0 ,
then $\gamma_{ab*} : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is injective;
- (iii) If ab satisfies the $(p - 2)$ - and $(p - 1)$ -, and p -link conditions in L_0 ,
then $\gamma_{ab*} : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is an isomorphism.

Proof. We prove only (i) from which the proofs for (ii) and (iii) become obvious. If ab satisfies the i -link conditions for $i = p, p - 2, p - 1$ in L_0 , then it satisfies the same conditions for L as well. Now apply Theorem 3.1 from previous section to conclude the following:

$$\begin{aligned} \gamma_{ab*} : H_p(L) &\rightarrow H_p(L') \text{ is surjective if } ab \text{ satisfies } (p - 1)\text{-link condition in } L_0 \\ \gamma_{ab*} : H_{p-1}(L_0) &\rightarrow H_{p-1}(L'_0) \text{ is surjective if } ab \text{ satisfies } (p - 2)\text{-link condition in } L_0 \\ \gamma_{ab*} : H_{p-1}(L) &\rightarrow H_{p-1}(L') \text{ is injective if } ab \text{ satisfies } (p - 1)\text{-link condition in } L_0 \end{aligned}$$

Now apply Theorem 4.1(i) to finish the proof of (i). □

It is important to notice that the required link conditions in Theorem 4.2 need to be satisfied in $L_0 \subset L \subseteq K$. It is not true that if ab satisfies the i -link condition in K , then it does so in L and L_0 . However, for $i = p - 1, p - 2$, we have the following observation which allows us to extend the results in Theorem 4.2(i) to the case when ab satisfies $(p - 1)$ - and $(p - 2)$ -link conditions in K .

Proposition 4.3. *Let $\gamma_{ab} : K \rightarrow K'$. Let (L', L'_0) be any pair of subcomplexes of dimensions $p + 1$ and p , respectively, where $L' \subseteq K'$ and $L'_0 \subset L'$. Let L' contain the vertex a . There exists a pair of subcomplexes (L, L_0) of dimensions $p + 1$ and p , respectively, where $L \subseteq K$ and $L_0 \subset L$ so that $\gamma_{ab} : (L, L_0) \rightarrow (L', L'_0)$ and, for $i = p - 1, p - 2$, ab satisfies the i -link condition for L_0 if it does so for K .*

Proof. Consider the preimage of L' under γ_{ab} and take L as its $(p + 1)$ -dimensional skeleton. Similarly take L_0 as the p -dimensional skeleton of the preimage of L'_0 . For $i = p - 1, p - 2$, if there is a simplex σ of dimension $(i - 1)$, in $\text{Lk } a \cap \text{Lk } b$ in L_0 , then the simplex $ab * \sigma$ has to be present in L_0 since this simplex is in K because ab satisfies the i -link condition in K . Therefore, ab satisfies the same link condition in L_0 as well. □

Notice that the above observation does not include the p -link condition. Since the p -link condition may require a $(p + 1)$ -simplex, the edge ab may not satisfy the p -link condition in the p -dimensional complex L_0 , even if it does so in K . This is why we cannot extend Theorem 4.2(ii) and (iii) in the result below which relates the surjectivity of the relative homology groups with the link conditions in the original input complex K instead of a subcomplex. The proof of this theorem follows from applying Proposition 4.3 to Theorem 4.2(i).

Theorem 4.4. *Let $L' \subseteq K'$ be any pure subcomplex of dimension $p + 1$ and $L'_0 \subset L'$ be any of its pure subcomplex of dimension p , where $K' = \gamma_{ab}(K)$. There exists a pair of pure subcomplexes L, L_0 of dimensions $p + 1$ and p , respectively with $L \subseteq K$ and $L_0 \subset L$ so that:*

If ab satisfies $(p - 2)$ - and $(p - 1)$ -link conditions in K , then $\gamma_{ab} : H_p(L, L_0) \rightarrow H_p(L', L'_0)$ is surjective.*

4.1 Implications of Theorem 4.4 and Theorem 4.2

Surjectivity. Notice that Theorem 4.4 is sufficient to claim that edge contractions cannot create any new class in relative homology groups if the edge ab satisfies only $(p - 1)$ - and $(p - 2)$ -link conditions in K . However, it may happen that the surjectivity does not respect torsions, that is, the preimage of a torsion

subgroup in $H_p(L', L'_0)$ may not have torsion in $H_p(L, L_0)$. This would mean that $(p - 1)$ - and $(p - 2)$ -link conditions are not enough to guarantee that torsion subgroups have sujection only from the torsion subgroups of the source space. We are interested in link conditions that guarantee that no new torsion class is created by an edge contraction because of its connection to the problem of OHCP as discussed in the next section.

We present an example where a new relative torsion indeed appears in the contracted complex even though the edge satisfies the $(p - 1)$ - and $(p - 2)$ -link conditions. Consider a sequence of triangles forming a Möbius strip. It is known that Möbius strip has a torsion in H_1 relative to its boundary. Now remove one triangle, say abc , and call the new complex M . Assume that the edge ab was on the boundary of M . The complex M , which is a Möbius strip with one triangle removed, does not have any relative torsion. However, if we contract the edge ab , we eliminate the hole created by the removal of abc . The resulting complex now is a Möbius strip and hence has a relative torsion. Notice that, for $p = 1$, ab satisfied $(p - 1)$ - and $(p - 2)$ -link conditions vacuously though it did not satisfy the p -link condition.

The isomorphism result in Theorem 4.2(iii) implies that no *new* relative torsion appears if ab satisfies all three link conditions, namely p -, $(p - 1)$ -, and $(p - 2)$ -link conditions in L_0 , which cannot be guaranteed by requiring ab to satisfy them in K .

In section 6, by a graph theoretic approach, we show that the p -link condition in K alone is sufficient to prevent the appearance of any *new* relative torsion. The above example illustrates that the p -link condition is sometimes necessary for preventing new relative torsions from appearing. Notice that we are not concerned with the relative torsions existing before the edge contraction.

Injectivity. Unlike surjection, the injectivity implied by Theorem 4.2(ii) cannot be used to claim that no relative homology class will be killed by an edge ab satisfying the p - and $(p - 1)$ -link conditions in K . The reason is that for injectivity of γ_{ab*} we choose the pair (L, L_0) in K first and then consider its image under γ_{ab} . So, even if ab satisfies p - and $(p - 1)$ -link conditions in K , it may not do so in L_0 and hence Theorem 4.2(ii) may not be applied. In fact, contracting edges satisfying all link conditions in K can indeed kill a relative torsion. Figure 3 illustrates such an example. The resulting complex after contracting edge ab in the 2-complex on left is shown on right. The left complex minus the triangles abd , abm and the edge ab forms a 15-triangle Möbius strip with a self-intersection at vertex d . This strip relative to its boundary results in a relative torsion. On the right, the self intersection expands to the *edge* ad , which causes the relative torsion to disappear.

This example illustrates that an edge contraction may destroy a relative torsion though no new relative torsion is generated thanks to Theorem 4.4(iii). This is of course good news for OHCP since such edge contractions can only take away obstructions to their solutions and not introduce new ones. We elaborate on this statement now, and describe a concrete example illustrating the benefits of such edge contractions to the efficient solution of OHCP instances.

4.2 Relative homology preservation and the OHCP

Given an oriented simplicial complex K of dimension d , and a natural number p , $1 \leq p \leq d$, the p -boundary matrix of K , denoted $[\partial_p]$, is a matrix containing exactly one column j for each p -simplex σ in K , and exactly one row i for each $(p - 1)$ -simplex τ in K . If τ is not a face of σ , then the entry in row i and column j is 0. If τ is a face of σ , which we denote by $\tau \preceq \sigma$, then this entry is 1 if the orientation of τ agrees with the orientation induced by σ on τ , and -1 otherwise.

Note that given a simplicial complex K and a choice of orientations for all simplices in K , $[\partial_p]$ is unique under row and column permutations, and will generally be referenced as a single matrix. We also need the notion of total unimodularity. A matrix A is *totally unimodular*, or *TU*, if the determinant of each square submatrix of A is either 0, 1, or -1 . An immediately necessary condition for A to be TU is that each




Figure 3: Contracting ab destroys a relative homology class even though it satisfies the link condition. In the left complex, a Möbius strip is formed by the triangles with orientations shown in black. This Möbius strip does not exist in the right complex, since it self-intersects along edge ad here. The input chain for the OHCP instance is edge ef (shown in purple). Candidates for optimal homologous chains are shown in green and cyan on the left complex, and in green on the right complex. See Section 4.3 for more details.

$a_{ij} \in \{0, \pm 1\}$. The importance of TU matrices for integer programming is well known – see, for instance, the book by Schrijver [20, Chapters 19-21]. In particular, it is known that the *integer* linear program

$$\min \mathbf{f}^T \mathbf{x} \quad \text{subject to} \quad A\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0} \text{ and } \mathbf{x} \in \mathbb{Z}^n$$

for $A \in \mathbb{Z}^{m \times n}$, $\mathbf{b} \in \mathbb{Z}^m$ can *always*, i.e., for every $\mathbf{f} \in \mathbb{R}^m$, be solved in polynomial time by solving its linear programming *relaxation* (obtained by ignoring $\mathbf{x} \in \mathbb{Z}^n$) *if and only if* A is *totally unimodular*. The main motivation for our discussion of totally unimodular matrices is the following result in [9].

Theorem 4.5. *For a finite simplicial complex K of dimension greater than p , the boundary matrix $[\partial_{p+1}]$ is totally unimodular if and only if $H_p(L, L_0)$ is torsion-free, for all pure subcomplexes L_0, L in K of dimensions p and $p + 1$ respectively, where $L_0 \subset L$.*

Following this result, instances of OHCP and related problems (with homology over \mathbb{Z}) can be solved in polynomial time when the simplicial complex K is relative torsion-free. It is appropriate to mention here that the connection between total unimodularity of boundary matrices and torsion in the complex had been observed as early as 1895 by Poincaré [18] (described in his own language). However, the result in [9] connecting the total unimodularity with relative torsion is different, and has led to a polynomial time algorithm for the OHCP problem. Notice that a complex can be torsion-free, but have nontrivial relative torsion. The Möbius strip is such an example.

4.3 An example where edge contraction helps to solve the OHCP

We present a small example which illustrates the effectiveness of edge contractions on efficient solutions of OHCP. Consider the simplicial complexes K on the left and K' on the right in Figure 3. We obtain K' from K by contracting the edge ab , which satisfies the link condition.

We consider the following OHCP instance on K , and equivalently on K' . All edges in red and the edge ef in purple have weights of 1 each. All edges in green and the thinner edges in black have weights of 0.05 each. The two edges bd and bm in K , drawn in cyan, have weights of 0.10 each. Orientations of all triangles and the edge ef are shown. Remaining edges could be oriented arbitrarily. The input 1-chain \mathbf{c} consists of the single edge ef with multiplier 1.

The unique optimal solution to the OHCP LP is the chain \mathbf{x} consisting of the 15 edges shown in green and cyan in K , with multipliers of ± 0.5 , resulting in a total weight of 0.425. \mathbf{x} is homologous to \mathbf{c} over \mathbb{R} , as their difference is the boundary of the 2-chain with multipliers of -0.5 for each of the 15 triangles in the Möbius strip, whose orientations are shown using black arrows (triangles abd and abm have multipliers of 0). But the unique optimal homologous chain sought here is the 1-chain \mathbf{x}' shown in green in the right complex (this chain is identical in K and K'). Consisting of the 9 edges $ak, km, ma, ec, en, nl, lj, jd, df$ with multipliers ± 1 each, \mathbf{x}' has a total weight of 0.45. But to obtain this solution, we have to solve the OHCP model as an *integer* program, instead of as a linear program.

Notice that K' obtained from K by contracting edge ab is free of relative torsion, and the weights of each edge in K' is the same as its weight in K . The unique optimal solution to the OHCP LP on K' for the identical input chain is \mathbf{x}' itself. \mathbf{x}' is homologous to \mathbf{c} as defined by the 2-chain consisting of the triangles whose orientations are shown using blue arrows, each with a multiplier of -1 . Hence we are able to solve the original OHCP instance using linear rather than integer programming, after contracting one edge.

We show that satisfying the p -link condition alone is sufficient to guarantee that no new relative torsion is introduced. Notice that this result says nothing about any existing relative torsion. Instead of proving the result directly, we show that the p -link condition for the edge being contracted preserves the total unimodularity of the boundary matrix, which in turn guarantees the absence of relative torsion thanks to Theorem 4.5. In contrast to our earlier approach, we use results from graph theory to arrive at this result.

5 Bipartite Graphs and Boundary Matrices

Given a matrix A with entries in $\{0, \pm 1\}$, we associate with it a weighted, undirected, bipartite graph $G = (V_1 \cup V_2, E)$ where each row of A corresponds to a node in V_1 and each column of A corresponds to a node in V_2 [4, 6]. Each nonzero entry a_{ij} in A is associated with an edge connecting the nodes in V_1 and V_2 corresponding to row i and column j , respectively, with a weight equal to a_{ij} . We call G the bipartite graph representation of A .

Some definitions from graph theory are central to our discussion. For a subgraph S of G containing a vertex v , we denote the number of edges in S incident to v as the *degree* of v in S , or $\deg_S(v)$. A *cycle* Y is a connected subgraph of G where for each vertex $v \in Y$, $\deg_Y(v) = 2$. We call a subgraph C of G a *circuit* if for each vertex $v \in C$, $\deg_C(v) = 0 \pmod{2}$. By this definition, it is possible that $v \in C$, but $\deg_C(v) = 0$. However, we consider two circuits as equivalent if they contain the same edge sets with the same weights for each edge. This means, for example, that a circuit C containing only vertices of degree 0 in C is equivalent to the empty subgraph.

Definition 5.1. A circuit C in the weighted graph G is *b-even* if the sum of the weights of the edges in C is $0 \pmod{4}$, and *b-odd* if the sum of the weights of the edges is $2 \pmod{4}$. The quality of C being b-even, b-odd, or neither is called the *b-parity* of C .

This definition is equivalent to the definition of *even* and *odd* cycles given by Conforti and Rao [6]. Given a simplicial complex K , consider the p -graph $G_p(K)$ of K constructed as follows. Each p - and $(p-1)$ -simplex $\sigma \in K$ provides a dual vertex σ^* in $G_p(K)$. We call the vertex σ^* a p - or $(p-1)$ -vertex if σ is a p - or $(p-1)$ -simplex, respectively. There is an edge $\sigma^* \tau^*$ in $G_p(K)$ if and only if the $(p-1)$ -simplex $\sigma \in K$ has the p -simplex τ as a coface. The weight of $\sigma^* \tau^*$ is 1 or -1 depending on whether σ and τ match in orientation or not, respectively. It is evident that $G_p(K)$ is a weighted bipartite graph whose adjacency matrix is given by $[\partial_p]$. This means the following proposition is almost immediate.

Proposition 5.2. *Reversing the orientation of any collection of p - and $(p-1)$ -simplices in K does not alter the b -parity of any circuit in $G_p(K)$.*

Let C be a circuit in a graph G . A *chord* of C is a single edge of G not in C whose both end points are vertices in C . If C has no chord, it is called *chordless*, and to say C is induced is equivalent to saying C is chordless. Using the terminology given above, we now state without proof an important result presented by Conforti and Rao [6], who were in turn motivated by the results of Camion [4].

Theorem 5.3. *For a matrix A with entries in $\{0, \pm 1\}$ and its bipartite graph representation G , A is totally unimodular if and only if G contains no chordless b -odd circuit.*

By combining Theorem 4.5 and Theorem 5.3, the following corollary is immediate.

Corollary 5.4. *For a finite simplicial complex K of dimension greater than p , the following results are equivalent.*

1. $H_p(L, L_0)$ is torsion-free for all pure subcomplexes L_0, L in K of dimensions p and $p+1$, respectively, where $L_0 \subset L$.
2. The boundary matrix $[\partial_{p+1}]$ is totally unimodular.
3. The bipartite graph representation $G_{p+1}(K)$ of $[\partial_{p+1}]$ contains no chordless b -odd circuit.

5.1 Simplices and edge contraction

Recall that each simplex in K is mapped by an edge contraction $\gamma_{ab} : K \rightarrow K'$ to a simplex in K' . We categorize simplices into three cases based on how they get mapped by γ_{ab} . These cases are defined relative to the specific edge ab being contracted. We illustrate these cases in Figure 4, and introduce several related definitions below.

1. For each pair of simplices $\sigma, \sigma' \in K$ where $a \in \sigma, b \in \sigma'$, and $\sigma = (\sigma' \setminus \{b\}) \cup \{a\}$, we have $\gamma_{ab}(\sigma) = \gamma_{ab}(\sigma') = \sigma$. Then σ and σ' are *mirror* simplices, and we say they are mirrors of each other. Their duals are mirror vertices, and we say these vertices are mirrors of each other.

In Figure 4, a is the mirror of b , ad is the mirror of bd , ae is the mirror of be , and ade is the mirror of bde . Similarly, dual vertex u_1 is the mirror of u_2 , v_1 is the mirror of v_2 , and w_1 is the mirror of w_2 .

2. $\sigma \in K$ is *collapsing* if $a, b \in \sigma$. Its dual σ^* is a *collapsing* vertex. Note if σ is p -dimensional, then σ has exactly one pair of $(p-1)$ -faces that are mirrors of each other. Then $\gamma_{ab}(\sigma) = \tau$, where τ is the unique mirror $(p-1)$ -face of σ containing a .

In Figure 4, ab, abd, abe , and $abde$ are all collapsing simplices. v_3, w_3 , and s are collapsing vertices.

3. $\sigma \in K$ is *injective* if neither of the above cases applies. We have $\gamma_{ab}(\sigma) = \sigma'$, and $\gamma_{ab}^{-1}(\sigma') = \sigma$. If $b \notin \sigma$, then $\sigma' = \sigma$. If $b \in \sigma$, then $\sigma' = (\sigma \setminus \{b\}) \cup \{a\}$.




Figure 4: The 2-graph of a complex, and the result after edge contraction γ_{ab} .

We state a few more definitions related to the Mirror and Collapsing cases.

Definition 5.5. Let $\sigma_1 \in K$ be a $(p+1)$ -simplex with mirror $\sigma_2 \in K$. Note that σ_1 and σ_2 have exactly one common p -face ξ that is an injective simplex. Each other p -face $\tau_1 \preceq \sigma_1$ is a mirror of some $\tau_2 \preceq \sigma_2$, and vice-versa. In any subgraph S of the dual graph G_{p+1} of K that contains both edges $\sigma_1^* \tau_1^*$ and $\sigma_2^* \tau_2^*$, these two edges are *mirror edges*, and are mirrors of each other. In Figure 4, assuming S is the entire 2-graph, tu_1 is the mirror of tu_2 , $u_1 v_1$ is the mirror of $u_2 v_2$, and $u_1 w_1$ is the mirror of $u_2 w_2$.

Definition 5.6. For any p -simplices τ_1 and τ_2 that are mirrors of each other in K , the two unique edges of the dual graph G_{p+1} of K that directly connect τ_1 and τ_2 each to their common collapsing $(p+1)$ -coface σ are together called the *mirror connection* between τ_1 and τ_2 , and also between τ_1^* and τ_2^* . There are two mirror connections in Figure 4. The mirror connection between v_1 and v_2 are the two edges $v_3 v_1$ and $v_3 v_2$, and the mirror connection between w_1 and w_2 are $w_3 w_1$ and $w_3 w_2$.

Definition 5.7. Any edges incident to two vertices in the dual graph G_{p+1} that are both collapsing are called *collapsing edges*. In Figure 4, sv_3 and sw_3 are collapsing edges.

6 Relative Torsion-Aware Edge Contractions

We now present the theorem which states that contracting an edge in a simplicial complex satisfying the p -link condition does not create any new relative torsion in dimension p . In other words, if the $(p+1)$ -boundary matrix of a simplicial complex is totally unimodular to start with, such an edge contraction will preserve this property. Theorem 6.1 below states the same thanks to Corollary 5.4.

Theorem 6.1. Let $K' = \gamma_{ab}(K)$ where ab satisfies the p -link condition. Let \mathcal{C} be the set of circuits in graph $G = G_{p+1}(K)$, and \mathcal{C}' be the set of circuits in $G' = G_{p+1}(K')$. Let $\mathcal{C}_L \subseteq \mathcal{C}$ be the set of all circuits of G that contain a collapsing p -vertex. Let $\mathcal{C}_M \subseteq \mathcal{C}$ be the set of all circuits of G that contain a pair of $(p+1)$ -vertices that are mirrors of each other. There exists a function $f : \mathcal{C} \setminus (\mathcal{C}_M \cup \mathcal{C}_L) \rightarrow \mathcal{C}'$ with the following properties.

- P1.** f is surjective.
- P2.** f preserves b -parity.
- P3.** For each C in the domain of f , if C has a chord then so does $f(C)$.

Proof. Before providing details of the proof, we refer the reader to Figure 4 to visualize \mathcal{C}_M and \mathcal{C}_L for $p = 1$. \mathcal{C}_L is the set of all circuits that contain edges incident to the vertex s , which is the dual of the only collapsing p -simplex in this complex. One such circuit is $\{sv_3, v_3v_1, v_1u_1, u_1w_1, w_1w_3, w_3s\}$. \mathcal{C}_M is the set of circuits which contain an edge incident to u_1 , and also an edge incident to u_2 . One such circuit is $\{u_1t, tu_2, u_2v_2, v_2v_3, v_3v_1, v_1u_1\}$. Though there are several circuits in this graph, none are in the domain of f as all of them belong to either \mathcal{C}_M or \mathcal{C}_L .

We break the proof up into several parts. We first define the function f in Section 6.1, and then show that f is surjective (Property **P1**) in Section 6.1.1. Finally, we prove that f preserves b -parity (Property **P2**) in Section 6.1.2 and that f preserves chords of circuits (Property **P3**) in Section 6.1.3.

6.1 Definition of f

We define the function f in the following way. For any $C \in \text{Dom}(f)$, let $L = \{\sigma \mid \sigma^* \in C\}$. For each vertex σ^* of C , $f(\sigma^*) = \gamma_{ab}(\sigma)^*$. The function f maps each edge $\sigma^*\tau^* \in C$ to the edge $\gamma_{ab}(\sigma)^*\gamma_{ab}(\tau)^*$ if the edge is not a mirror connection, otherwise simply to the vertex $\gamma_{ab}(\sigma)^* = \gamma_{ab}(\tau)^*$. Since γ_{ab} is well defined, for any $C \in \text{Dom}(f)$, if $f(C) = X$ and $f(C) = Y$, then $X = Y$.

To show $f(C) \in \mathcal{C}'$, first note that if all simplices of L are injective, this is almost trivially so since in these cases, γ_{ab} is injective, preserves the dimension of simplices, and $\tau \in L$ is a face of $\sigma \in L$ if and only if $\gamma_{ab}(\tau)$ is a face of $\gamma_{ab}(\sigma)$.

If L does not contain any collapsing simplices, but contains a pair of mirror simplices τ_1 and τ_2 , then since $C \notin \mathcal{C}_M$, τ_1 and τ_2 are p -simplices. Furthermore, the proper cofaces of either τ_1 or τ_2 are injective because the only coface which cannot be such is a collapsing simplex. Let $\tau = \gamma_{ab}(\tau_1) = \gamma_{ab}(\tau_2)$. Since γ_{ab} maps the proper cofaces of τ_1 and τ_2 injectively, $\deg_{f(C)}(\tau^*) = \deg_C(\tau_1^*) + \deg_C(\tau_2^*)$. Since both operands of this sum are even, the sum is even. Therefore, $f(C) \in \mathcal{C}'$.

Now assume that L contains a collapsing simplex σ . Since $C \notin \mathcal{C}_L$, σ is a $(p+1)$ -simplex. Furthermore, since all but two p -faces τ_1 and τ_2 of σ are also collapsing, σ^* must be of degree 2 in C . Therefore, C must contain the mirror connections $\tau_1^*\sigma^*$ and $\tau_2^*\sigma^*$. Note that σ is the only $(p+1)$ -simplex that is a coface of both τ_1 and τ_2 . The edge contraction γ_{ab} maps other proper cofaces of τ_1 and τ_2 in L injectively. Both mirror connections $\tau_1^*\sigma^*$ and $\tau_2^*\sigma^*$ are mapped by f to the vertex τ^* where $\tau = \gamma_{ab}(\tau_1) = \gamma_{ab}(\tau_2) = \gamma_{ab}(\sigma)$. Therefore, $\deg_{f(C)}(\tau^*) = \deg_C(\tau_1^*) + \deg_C(\tau_2^*) - 2$. Again, since $\deg_C(\tau_1^*) + \deg_C(\tau_2^*)$ is even, $\deg_{f(C)}(\tau^*)$ is even. Hence, $f(C) \in \mathcal{C}'$.

6.1.1 Surjectivity of f

We have shown that f maps each element in its domain to an element in its codomain. To show that f is surjective (Property **P1**), we define a function that extends the domain of f to all subgraphs of G . Let \mathcal{S} be the set of all subgraphs of G , and \mathcal{S}' be the set of all subgraphs of G' . Define the function $g : \mathcal{S} \rightarrow \mathcal{S}'$ as the extension of f to this domain. Hence it is possible that some vertex $\tau^* \in \text{Dom}(g)$ is the dual of a collapsing p -simplex τ . In this case we specify that g maps τ^* and all edges of G incident to τ^* , which all must be collapsing edges, to the empty subgraph. Since γ_{ab} is surjective, so is g . To show f is surjective, we will take an arbitrary $C' \in \mathcal{C}'$ and show that its preimage under g , which we will denote \mathcal{S}_C , must contain an element of $\text{Dom}(f)$.

If \mathcal{S}_C contains an element $S \notin \text{Dom}(f)$, there are three type of edges that we may remove or add to S without affecting $g(S)$, as detailed below.

- We may remove or add any collapsing edge.
- We may remove or add any edge that is a mirror connection in G .

- We may remove either one, but not both, of any two edges that are mirrors of each other. For any edge in S that is a mirror in G but not in S , we may add its mirror to S .

Using these modifications, we describe a procedure to construct a $C_f \in \text{Dom}(f) \cap \mathcal{S}_C$ (see Figure 5). Figures 6, 7, and 8 illustrate the steps of this procedure. Graph vertices are labeled with the vertex labels of the dual simplices. In each figure, graph edges and vertices eliminated by the step shown are highlighted in red, and edges and vertices added by the step are highlighted in green.

PROCEDURE CONSTRUCT C_f
 Input: $S \in \mathcal{S}_C \setminus \text{Dom}(f)$ in graph G such that $g(S) \in \mathcal{C}'$.
 Output: $C_f \in \mathcal{S}_C \cap \text{Dom}(f)$ such that $g(C_f) = g(S)$.

- I. Remove any collapsing edges.
- II. For any edge $\sigma^*\tau^* \in S$ that is a mirror in G , where σ adjoins the contracting vertex b , replace it with its mirror in G if this mirror is currently absent in S , or remove $\sigma^*\tau^*$ otherwise.
- III. For each pair of vertices of odd degree that are mirrors of each other, negate the mirror connections connecting them, i.e., remove them if they are in the current subgraph, or add them if they are not.
- IV. For each pair of vertices of odd degree incident to a common single edge in G , negate this edge.

Figure 5: Procedure to construct a circuit in $\text{Dom}(f)$ starting with a subgraph not in $\text{Dom}(f)$.




Figure 6: A subgraph S of the 3-graph where $g(S) = C' \in \mathcal{C}'$ (left) and changes in Step I (right).

We provide some details of the steps in the procedure now. Denote the subgraph after Step II as S_2 . Since we have removed any collapsing edges, the only edges remaining in S_2 that g does not map injectively are mirror connections in G .

Since C' is a circuit, the only vertices that are of odd degree in S_2 are vertices incident to a mirror connection in G .

For any pair of mirror connections connecting two mirror vertices in G , there are three vertices involved – two p -vertices τ_1^* and τ_2^* that are mirrors of each other in G , and a collapsing $(p+1)$ -vertex σ^* . $\deg_{S_2}(\sigma^*)$ is the number of edges of the mirror connection in S_2 , and these are the edges g removes by mapping them to the vertex $v = g(\tau_1^*) = g(\tau_2^*) = g(\sigma^*)$. Because C' is a circuit, and edges that are not a mirror connection are mapped injectively by g , $\deg_{S_2}(\sigma^*)$ is odd if and only if $\deg_{S_2}(\tau_1^*) + \deg_{S_2}(\tau_2^*)$ is odd. Therefore, among the three vertices τ_1^* , τ_2^* , σ^* , an even number of them must be of odd degree in S_2 . The graph on the right in Figure 7 illustrates this situation. Note that the vertices ahi^* and bhi^* have degrees 3 and 1 in S_2 , respectively, and the vertex $abhi^*$ has degree 2 in S_2 .




Figure 7: Changes in Step II (left) and in Step III (right).

If $\tau_1^*, \tau_2^*, \sigma^*$ are all of even degree, no further action is needed. If this is not the case, but σ^* is of even degree, then Step III will make τ_1^* and τ_2^* of even degree, and σ^* will still be of even degree. Notice that since ab satisfies the p -link condition, the mirror connections $\tau_1^*\sigma^*$ and $\tau_2^*\sigma^*$ must exist in G . If σ^* is of odd degree, then Step IV will convert σ^* and whichever vertex between τ_1^* and τ_2^* is of odd degree to vertices of even degree.

Steps III and IV ensure that every vertex in the subgraph have even degrees. The resulting subgraph C_f is a circuit in \mathcal{S}_C that is not in \mathcal{C}_M because we have removed at least one of the edges adjoining vertices that are mirror to each other in Step II, and not in \mathcal{C}_L because we have removed all collapsing edges in Step I. Therefore, $C_f \in \text{Dom}(f) \cap \mathcal{S}_C$.

6.1.2 Preservation of b -Parity

To show f preserves b -parity (Property **P2**), let $C \in \text{Dom}(f)$, $C' = f(C)$, and $L = \{\sigma \mid \sigma^* \in C\}$. We need to analyze only those vertices σ^* of C where σ is collapsing, or is a mirror. Since $C \in \text{Dom}(f)$, we may restrict our attention to mirror p -vertices in C , or collapsing $(p+1)$ -vertices in C . In fact, we may further restrict our attention to edges incident to p -vertices τ_1^* and τ_2^* that are mirrors and in C , because $C \in \text{Dom}(f)$ implies any edges incident to a collapsing $(p+1)$ -vertex is also incident to a p -vertex that is a mirror in C .




Figure 8: Changes in Step IV, and C_f .

We will divide the possibilities into four cases based on two criteria – either the mirror connections connecting τ_1^* and τ_2^* are both in C or both not in C , and either the orientations of τ_1 and τ_2 are consistent, or they are not. Recall that if τ_1 and τ_2 have a unique common $(p-1)$ -face ξ , τ_1 and τ_2 are consistently oriented if they induce opposite orientations on ξ . We assume that unless otherwise specified, the orientations of simplices in K' are chosen so that for each edge e in C , the weights of e and $g(e)$ are the same.

If neither edge of the mirror connection is in C , f maps all edges of C injectively to edges of C' . If the orientations of τ_1 and τ_2 are not consistent, then we may choose the ordering of vertices of $\tau = \gamma_{ab}(\tau_1)$ that determines its orientation to be the same as the ordering of the vertices of τ_1 . Then each edge e of C' will have the same weight as $g^{-1}(e)$ in C . If the orientations of τ_1 and τ_2 are consistent, then we may still choose the order of vertices of τ in the same way as described above, but now each edge e incident to τ_2^* will have the opposite weight as $g(e)$. The sum of the weights of edges incident to τ_2^* and the sum of the weights of their corresponding edges in C' must differ by 2 for each such edge e . Since C is a circuit, there must be an even number of such edges. Therefore, this difference is $0 \pmod 4$, and hence the b -parity of C' is the same as that of C . We now discuss the remaining case when both mirror connections are in C and the orientations of τ_1 and τ_2 are not consistent. We illustrate both cases for $p = 1$ in Figure 9.

Consider the case when both mirror connections are in C and the orientations of τ_1 and τ_2 are not consistent. Notice that the orientation of their common coface σ must agree with the orientation of one of these mirrors and disagree with the orientation of the other mirror. Hence, the sum of the weights of the mirror connections must be 0. The function f maps all other edges of C injectively to edges of C' , and we may choose the orientation of τ as before. Therefore, the sum of the weights of edges in C' equals the sum of the weights of edges in C .

If the orientations of τ_1 and τ_2 are consistent, then the sum of the weights of the mirror connections must be $2 \pmod 4$. However, the number of edges mapped injectively by g that are incident to τ_2^* must be odd. Hence, the difference between the sum of weights of these edges and their corresponding edges in C' must also be $2 \pmod 4$. It follows that the overall difference between the weights of edges in C and C' is $0 \pmod 4$,




Figure 10: A chord h of cycle C , and $C' = f(C)$.

necessary for absolute homology groups preservation. For relative homology, the scenario becomes more subtle. The injectivity and hence isomorphism cannot be guaranteed with any link condition in the original complex as our example in Figure 3 shows. The link conditions in Theorem 4.4 are sufficient for surjectivity; but are they necessary? This remains open. For relative torsions, Theorem 6.1 and our example of punctured Möbius strip in Section 4.1 show that the p -link condition is sufficient and sometimes necessary for preventing new torsions, and our example in Figure 3 shows that no link condition can ensure preserving existing relative torsions.

An open question is whether there exist local conditions that also preserve the total unimodularity property of the complex. We know that total unimodularity cannot be destroyed by an edge contraction if the appropriate link condition holds. But, we do not know how to preserve the *absence* of total unimodularity with local conditions. In the Section 4.3, we discuss a concrete example where linear programming fails to provide an optimal homologous chain, but an edge contraction makes it amenable to such computation by eliminating relative torsion. An immediately relevant question is how to approximate the optimal chain in the original complex using the one in a reduced complex.

Acknowledgment

We acknowledge the support of the NSF grants CCF-1064416, CCF-1064429, and CCF-1064600.

References

- [1] Dominique Attali, André Lieutier, and David Salinas. Efficient data structure for representing and simplifying simplicial complexes in high dimensions. In *Proceedings of the 27th annual ACM symposium on Computational geometry*, SoCG '11, pages 501–509, New York, NY, USA, 2011. ACM.
- [2] Dominique Attali, André Lieutier, and David Salinas. Collapsing Rips complexes. Manuscript, 2013.
- [3] Silvia Biasotti, Bianca Falcidieno, Daniela Giorgi, and Michela Spagnuolo. The hitchhiker’s guide to the galaxy of mathematical tools for shape analysis. *SIGGRAPH course notes*, 2012.
- [4] Paul Camion. Characterization of Totally Unimodular Matrices. *Proceedings of the American Mathematical Society*, 16(5):1068–1073, 1965.

- [5] Chao Chen and Daniel Freedman. Measuring and computing natural generators for homology groups. *Computational Geometry*, 43(2):169–181, 2010. Special Issue on the 24th European Workshop on Computational Geometry (EuroCG’08).
- [6] Michele Conforti and Mendu R. Rao. Structural properties and recognition of restricted and strongly unimodular matrices. *Mathematical Programming*, 38:17–27, 1987.
- [7] Vin de Silva and Robert Ghrist. Homological sensor networks. *Notices of the American Mathematical Society*, 54(1):10–17, 2007.
- [8] Tamal K. Dey, Herbert Edelsbrunner, Sumanta Guha, and Dmitry V. Nekhayev. Topology preserving edge contraction. *Publications de l’ Institut Mathematique (Beograd)*, 60:23–45, 1999.
- [9] Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM Journal on Computing*, 40(4):1026–1040, 2011. arxiv:1001.0338.
- [10] Nathan M. Dunfield and Anil N. Hirani. The least spanning area of a knot and the optimal bounding chain problem. In *Proceedings of the 27th ACM Annual Symposium on Computational Geometry*, SoCG ’11, pages 135–144, 2011.
- [11] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’97, pages 209–216, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [12] Robert Ghrist. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45:61–75, 2008.
- [13] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [14] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Mesh optimization. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’93, pages 19–26, New York, NY, USA, 1993. ACM.
- [15] Sharif Ibrahim, Bala Krishnamoorthy, and Kevin R. Vixie. Simplicial flat norm with scale. *Journal of Computational Geometry*, 2011. Under review; arxiv:1105.5104.
- [16] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. The Johns Hopkins University Press, Baltimore, USA, 2001.
- [17] James R. Munkres. *Elements of Algebraic Topology*. Addison–Wesley Publishing Company, Menlo Park, 1984.
- [18] Henri Poincaré. *Papers on Topology: Analysis Situs and Its Five Supplements*. History of Mathematics. American Mathematical Society, 2010.
- [19] Joseph J. Rotman. *An Introduction to Algebraic Topology*. Springer Verlag, New York, 1988.
- [20] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons Ltd., Chichester, 1986.
- [21] Alireza Tahbaz-Salehi and Ali Jadbabaie. Distributed coverage verification algorithms in sensor networks without location information. *IEEE Transactions on Automatic Control*, 55(8):1837–1849, 2010.

- [22] David W. Walkup. The lower bound conjecture for 3- and 4-manifolds. *Acta Mathematica*, 125:75–107, 1970.
- [23] Thomas Wolle and Hans L. Bodlaender. A note on edge contraction. Technical Report UU-CS-2004-028, Department of Information and Computing Sciences, Utrecht University, 2004.
- [24] Zoë Wood, Hugues Hoppe, Mathieu Desbrun, and Peter Schröder. Removing excess topology from isosurfaces. *ACM Transactions on Graphics*, 23(2):190–208, April 2004.