Topology of matching complexes of complete graphs via discrete Morse theory

Anupam Mondal   Sajal Mukherjee   Kuldeep Saha

Institute for Advancing Intelligence (IAI), TCG CREST, Kolkata, India
Academy of Scientific & Innovative Research (AcSIR), Ghaziabad, India

Bouc (1992) first studied the topological properties of M_n, the matching complex of the complete graph of order n, in connection with Brown complexes and Quillen complexes. Björner et al. (1994) showed that M_n is homotopically $(\nu_n - 1)$-connected, where $\nu_n = \lfloor \frac{n+1}{3} \rfloor - 1$, and conjectured that this connectivity bound is sharp. Shareshian and Wachs (2007) settled the conjecture by inductively showing that the ν_n-dimensional homology group of M_n is nontrivial, with Bouc's calculation of $H_1(M_7)$ serving as the pivotal base step. In general, the topology of M_n is not very well-understood, even for a small n. In the present article, we look into the topology of M_n, and M_7 in particular, in the light of discrete Morse theory as developed by Forman (1998). We first construct a gradient vector field on M_n (for $n \geq 5$) that doesn't admit any critical simplices of dimension up to $\nu_n - 1$, except one unavoidable 0-simplex, which also leads to the aforementioned $(\nu_n - 1)$-connectedness of M_n in a purely combinatorial way. However, for an efficient homology computation by discrete Morse theoretic techniques, we are required to work with a gradient vector field that admits a low number of critical simplices, and also allows an efficient enumeration of gradient paths. An optimal gradient vector field is one with the least number of critical simplices, but the problem of finding an optimal gradient vector field, in general, is an NP-hard problem (even for 2-dimensional complexes). We improve the gradient vector field constructed on M_7 in particular to a much more efficient (near-optimal) one, and then with the help of this improved gradient vector field, compute the homology groups of M_7 in an efficient and algorithmic manner. We also augment this near-optimal gradient vector field to one that we conjecture to be optimal.

Keywords: discrete Morse theory, complete graph, matching, abstract simplicial complex, gradient vector field, Morse homology

1 Introduction

The collection of all *matchings* or independent edge sets (i.e., sets of edges without common endvertices) in a graph constitutes an abstract simplicial complex, called the *matching complex* of the graph. The matching complex of a graph G may also be realized as the *independence complex* of the *line graph* of G. In particular, we denote the matching complex of the complete graph of order n (see Björner et al. (1994);

Bouc (1992); Shareshian and Wachs (2007); Wachs (2003)) by M_n. Topological properties of M_n were first studied by Bouc (1992), in connection with *Brown complexes* (Brown (1974, 1975)) and *Quillen complexes* (Quillen (1978)). Björner et al. (1994) proved the following result regarding the homotopical connectivity of M_n.

Theorem 1.1. *For all n, the matching complex M_n is homotopically $(\nu_n - 1)$-connected, where $\nu_n = \lfloor \frac{n+1}{3} \rfloor - 1$.*

They also conjectured that the connectivity bounds of Theorem 1.1 are sharp. Shareshian and Wachs (2007) settled the conjecture by inductively showing that the ν_n-dimensional homology group of M_n is nontrivial (see also Wachs (2003)). The crucial base step was Bouc's "hand" calculation of the first homology group of M_7. Here we remark that the topology of M_n is not very well-understood even for a small n. For example, the first nontrivial reduced homology group of M_{14} is not known, but Jonsson (2009) showed that, rather surprisingly, M_{14} is the only matching complex of a complete graph, for which the ν_n-dimensional homology group has torsion other than 3-torsion (see also Shareshian and Wachs (2007)). Also, very little is known about the higher dimensional homology groups of M_n in general.

Forman (1998) developed discrete Morse theory as a combinatorial analogue of (smooth) Morse theory (see also Chari (2000); Forman (2002); Knudson (2015); Kozlov (2020); Scoville (2019)). Over the years, the theory has turned out to be immensely useful in diverse fields of theoretical and applied mathematics, and also in computer science. The central notion of this theory is that of a *discrete Morse function* defined on a finite (abstract) simplicial complex (or a (regular) CW complex). It helps us understand the topology of the complex through an efficient cell decomposition (i.e., one with fewer cells than in the original decomposition) of the complex. In practice, however, instead of such functions, we usually consider an equivalent and more useful notion of (discrete) *gradient vector fields* on the complex. The homotopy type of the complex is determined by only the simplices (or cells) that are *critical* with respect to an assigned gradient vector field. Discrete Morse theoretic techniques also help us compute the homology groups, Betti numbers, etc. of a complex in a computationally efficient way, provided one manages to construct a sufficiently "good" gradient vector field on it, in the first place.

In Section 3 of this article, we first prove the following by explicitly constructing a gradient vector field on M_n.

Theorem 1.2. *There is a gradient vector field on M_n (for $n \geq 5$) with respect to which there are no critical simplices of dimension up to $\nu_n - 1$, except one 0-simplex.*

Theorem 1.1 follows as a natural consequence of the above. Existence of a gradient vector field with the same property was previously shown by Shareshian and Wachs (2007) (Section 9), with the help of the fact that the ν_n-skeleton of M_n is *vertex decomposable*, as established by Athanasiadis (2004). However, the task of homology computation via discrete Morse theory relies on an efficient enumeration of all critical simplices and all possible *gradient paths* of a specific type. Thus, in order to compute the homology groups efficiently, we are required to work with a gradient vector field that admits a low number of critical simplices of each dimension to begin with. An *optimal gradient vector field* is one with the least number of critical simplices, but the problem of finding an optimal gradient vector field, in general, is an NP-hard problem (even for 2-dimensional complexes) (Eğecioğlu and Gonzalez (1996); Joswig and Pfetsch (2006); Lewiner et al. (2003a,b)). In this context, we mention that Adiprasito et al. (2017), Benedetti and Lutz (2014), Benedetti et al. (2023) described a scheme for searching optimal discrete gradient vector fields with a random heuristic, which turned out to be successful, even in some cases with a large input size.

An efficient gradient vector field reduces the task of computing homology to the computation of (discrete) *Morse homology* groups, which are homology groups of a relatively simpler chain complex. Bouc's computation of the first homology group of M_7, which, as mentioned before, is the pivotal base step to determine the connectivity of M_n in general, relies on some ingenious, but somewhat ad hoc tricks. In Section 4, we present an efficient and algorithmic computation of the homology groups of M_7. We extend our previously constructed gradient vector field on M_7 to a *near-optimal* one, and use it to compute the Morse homology groups of M_7 in Subsection 4.2, and obtain the following.

Theorem 1.3. *The nontrivial discrete Morse homology groups of M_7 are the following:*

$$H_0(M_7) = \mathbb{Z}, H_1(M_7) = \mathbb{Z}_3, \text{ and } H_2(M_7) = \mathbb{Z}^{20}.$$

In Subsection 4.3, we augment the near-optimal gradient vector field on M_7 even further and get the following.

Theorem 1.4. *There is a gradient vector field on M_7 with respect to which there are 22 critical 2-simplices, two critical 1-simplices, and one critical 0-simplex.*

It follows from Theorem 1.3 that with respect to any gradient vector field on M_7, there is at least one critical 0-simplex, and there is at least one critical 1-simplex. Also, there are at least 21 critical 2-simplices (as $H_1(M_7)$ has a torsion). This naturally raises the question whether these lower bounds are sharp. However, we believe these bounds are not attainable, and thus a gradient vector field on M_7 that satisfies the requirements of Theorem 1.4 is indeed an optimal one. We pose this as a conjecture (Conjecture 5.1) in the Conclusion section.

2 Preliminaries

2.1 Basics of combinatorics and graph theory

An *abstract simplicial complex* (or simply, a complex) is a (finite, nonempty) collection, say $\mathcal{K}$, of finite sets with the property that if $\sigma \in \mathcal{K}$ and $\tau \subseteq \sigma$, then $\tau \in \mathcal{K}$. We note that the empty set is always in $\mathcal{K}$. If $\sigma \in \mathcal{K}$, then σ is called a *simplex* or a *face* of $\mathcal{K}$. If the simplex σ is a set of cardinality $d + 1$, then the dimension of σ is d, and we call σ a d-dimensional simplex (or simply, a d-simplex). We denote a d-simplex σ by $\sigma^{(d)}$ whenever necessary. The dimension of a complex $\mathcal{K}$, denoted by $\dim(\mathcal{K})$, is the largest dimension of its faces. The *vertex set* of a complex $\mathcal{K}$ is defined as $V(\mathcal{K}) = \cup_{\sigma \in \mathcal{K}} \sigma$ (i.e., the collection of all elements in all the faces of $\mathcal{K}$). The elements of $V(\mathcal{K})$ are called the vertices of the complex $\mathcal{K}$.

Any complex has a unique (up to a homeomorphism) *geometric realization*. However, we don't differentiate between a complex and its geometric realization while discussing its topological properties; it should be understood from the context.

The *f-vector* of a complex $\mathcal{K}$ is the integer sequence $(f_0, \ldots, f_{\dim(\mathcal{K})})$, where f_i is the number of i-dimensional faces of $\mathcal{K}$. The *Euler characteristic* of $\mathcal{K}$, denoted by $\chi(\mathcal{K})$, is given by

$$\chi(\mathcal{K}) = \sum_{i=0}^{\dim(\mathcal{K})} (-1)^i f_i.$$

A (simple, finite, undirected) *graph G* is an ordered pair of (disjoint) finite sets $(V(G), E(G))$, where $E(G) \subseteq \{e \subseteq V(G) : |e| = 2\}$. The sets $V(G)$ and $E(G)$ are called the *vertex set* and the *edge set* of the

graph G, respectively. We call an element of $V(G)$ a *vertex* of G and an element of $E(G)$ an *edge* in G. If $e = \{x, y\}$ is an edge, then x and y are *endvertices* of the edge e. Also, we denote the edge $e = \{x, y\}$ by xy (or, yx) for the sake of brevity. We note that a graph can be viewed as an abstract simplicial complex of dimension 1 or less (with the vertices and the edges being considered as 0-simplices and 1-simplices, respectively), and vice versa.

A *complete graph* with n (≥ 1) vertices is a graph G with $|V(G)| = n$ and $E(G) = \{e \subseteq V(G) : |e| = 2\}$. This graph is unique (up to a graph isomorphism), and we denote it by K_n.

A *matching* α in a graph G is a set of edges with the property that if the edges e_1 and e_2 are in α, then $e_1 \cap e_2 = \emptyset$ (i.e., e_1 and e_2 don't share an endvertex). If the vertex x is an endvertex of an edge in a matching α, then we say α *covers* or *matches* x, and otherwise, x is *uncovered* or *unmatched* by α. Moreover, we say the matching α matches x with another vertex y if the edge xy is in α. A *perfect matching* is one that covers all the vertices of the graph.

We note that the collection of all matchings in a graph G is an abstract simplicial complex, and we call it the *matching complex* of G. In particular, for all n, we denote the matching complex of K_n by M_n. We may verify that for all n, $\dim(M_n) = \lfloor \frac{n}{2} \rfloor - 1$, and if $(f_0, \ldots, f_{\lfloor \frac{n}{2} \rfloor - 1})$ is the f-vector of M_n, then for all $i \in \{0, \ldots, \lfloor \frac{n}{2} \rfloor - 1\}$,

$$f_i = \frac{(i+1)!}{2^{i+1}} \binom{n}{2(i+1)} \binom{2(i+1)}{i+1}. \tag{2.1}$$

Example 2.1. Let $V(K_n) = \{1, \ldots, n\}$. We describe (the topology of) M_n, for all n up to 8.

1. $M_1 = \{\emptyset\}$.

2. $M_2 = \{\emptyset, \{12\}\}$, which is a single point.

3. $M_3 = \{\emptyset, \{12\}, \{13\}, \{23\}\}$, which is the space with three distinct points.

4. The maximal simplices of M_4 are $\{12, 34\}$, $\{13, 24\}$, and $\{14, 23\}$. Thus, M_4 is a space with three mutually disjoint 1-simplices, which is homotopy equivalent to the space with three distinct points.

5. As $\dim(M_5) = 1$, the complex M_5 can be viewed as a (connected) graph with the vertex set $\{ij : i, j \in [5], i < j\}$ and the edge set $\{\{ij, k\ell\} : i, j, k, \ell \in [5], i < j, k < \ell, \{i, j\} \cap \{k, \ell\} = \emptyset\}$. This is the well-known Petersen graph as shown in Fig. 1.

   Any connected graph G is homotopy equivalent to a wedge of k circles, where k is the cyclomatic number (also known as the circuit rank or cycle rank) of the graph, which is the number $|E(G)| - |V(G)| + 1$. Thus, M_5 is homotopy equivalent to the wedge of $15 - 10 + 1 = 6$ circles.

   Here we also recall that the cyclomatic number of a connected graph is the same as the first Betti number of the graph when considered as a simplicial complex.

6. Although $\dim(M_6) = 2$, we observe that each 1-simplex in M_6 is contained in exactly one 2-simplex. By the notion of (elementary) *collapses* in topology (Cohen, 1973, Chapter 1), the complex M_6 deformation-retracts to a complex of dimension 1. Thus, M_6 is also homotopy equivalent to a (connected) graph, and subsequently, to a wedge of k circles. Here we determine k from the Euler characteristic (which is a homotopy invariant) of M_6. If (f_0, f_1, f_2) is the f-vector of M_6, then from Equation (2.1), $f_0 = 15$, $f_1 = 45$, $f_2 = 15$, and thus $1 - k = \chi(M_6) = f_0 - f_1 + f_2 = -15$. Therefore, M_6 is homotopy equivalent to the wedge of 16 circles.
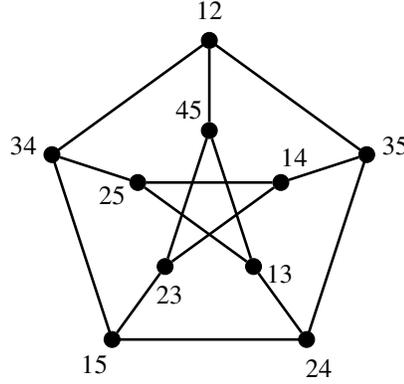
Fig. 1: The matching complex M_5 is the Petersen graph.

7. M_7 is a 2-dimensional complex, whose first homology group is not torsion-free (Bouc (1992)). The zeroth, first, and second homology groups of M_7 are $\mathbb{Z}$, $\mathbb{Z}_3$, and $\mathbb{Z}^{20}$, respectively (see Subsection 4.2).

8. M_8 is homotopy equivalent to the wedge of 132 spheres of dimension 2 (follows from Björner et al. (1994); see Example 3.3 below for a discrete Morse theoretic approach).

2.2 Simplicial homology

First, we need to introduce the notion of an orientation of a simplex. An orientation of a simplex is given by an ordering of its vertices, with two orderings defining the same orientation if and only if they differ by an even permutation. We denote an oriented k-simplex consisting of the vertices $x_0, x_1, \ldots, x_k$, with the orientation given by the increasing ordering of the indices, by $[x_0, x_1, \ldots, x_k]$. We usually choose and fix an ordering of the vertices of the complex to begin with, and assign each simplex the orientation corresponding to the induced ordering of its vertices. In other words, if $\sigma = \{x_0, x_1, \ldots, x_k\}$ is a k-simplex of a complex $\mathcal{K}$, and $x_0 < x_1 < \cdots < x_k$ with respect to the chosen order on $V(\mathcal{K})$, then in order to avoid notational complicacy, we denote the oriented k-simplex $[x_0, x_1, \ldots, x_k]$ by σ as well.

A k-chain in a complex $\mathcal{K}$ is a *finite* formal sum $\sum c_i \sigma_i$, where each $c_i \in \mathbb{Z}$ and each σ_i is an oriented k-simplex, with the notion that an oriented simplex is equal to the negative of the simplex with the opposite orientation (e.g., $[x_0, x_1, x_2, \ldots, x_k] = -[x_1, x_0, x_2, \ldots, x_k]$).

We denote the free abelian group generated by all k-simplices of a complex $\mathcal{K}$, i.e., the group of k-chains, by $C_k(\mathcal{K})$. We now define a homomorphism $\partial_k : C_k(\mathcal{K}) \to C_{k-1}(\mathcal{K})$ called the *boundary operator*. If $\sigma = [x_0, x_1, \ldots, x_k]$, considered as a basis element of $C_k(\mathcal{K})$, then

$$\partial_k(\sigma) := \sum_{i=0}^{k} (-1)^i [x_0, \ldots, \widehat{x_i}, \ldots, x_k],$$

where $[x_0, \ldots, \widehat{x_i}, \ldots, x_k]$ is the oriented $(k-1)$-simplex obtained from σ after deleting x_i (and with the induced orientation). We then extend ∂_k linearly to all k-chains. We also define ∂_0 to be the zero map.

In $C_k(\mathcal{K})$, the elements of the subgroup $\ker(\partial_k)$ are called *cycles* (more specifically, k-cycles), and the elements of the subgroup $\mathrm{im}(\partial_{k+1})$ are called *boundaries* (more specifically, k-boundaries).

We may verify that for all $k \geq 1$, and for any k-chain τ, we have $\partial_{k-1} \circ \partial_k(\tau) = 0$. In other words, $(C_*(\mathcal{K}), \partial_*)$ is a *chain complex*. The k-th homology group of $\mathcal{K}$, denoted by $H_k(\mathcal{K})$, is given by

$$H_k(\mathcal{K}) := {}^{\ker(\partial_k)} \big/ {}_{\mathrm{im}(\partial_{k+1})}.$$

The k-th *Betti number* of $\mathcal{K}$ is the free rank (i.e., rank of the torsion-free part) of $H_k(\mathcal{K})$.

We refer to the book *Elements of Algebraic Topology* by Munkres (1996) for background on algebraic and combinatorial topology, and the book *Graph Theory* by Diestel (2010) for background on graph theory.

2.3   Basics of discrete Morse theory

First, we introduce the notion of a *discrete vector field* and a (discrete) *gradient vector field* on an abstract simplicial complex following Forman (1998, 2002).

Definition 2.2 (Discrete vector field). A discrete vector field $\mathcal{V}$ on an abstract simplicial complex $\mathcal{K}$ is a collection of ordered pairs of simplices of the form (α, β) such that

  (i) $\alpha \subsetneq \beta$,

 (ii) dimension of the simplex β is 1 more than that of α,

(iii) each face of $\mathcal{K}$ is in *at most* one pair of $\mathcal{V}$.

If the simplex $\alpha^{(p)}$ is paired off with the simplex $\beta^{(p+1)}$ in a discrete vector field (i.e., the pair of simplices $(\alpha^{(p)}, \beta^{(p+1)})$ is an element of the discrete vector field), then we denote it by $\alpha \rightarrowtail \beta$ (or $\beta \leftarrowtail \alpha$).

Given a discrete vector field $\mathcal{V}$ on a simplicial complex $\mathcal{K}$, a $\mathcal{V}$-path is a sequence of simplices

$$\alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \ldots, \alpha_k^{(d)}, \beta_k^{(d+1)}, \alpha_{k+1}^{(d)}$$

such that for each $i \in \{0, \ldots, k\}$, $(\alpha_i, \beta_i) \in \mathcal{V}$ and $\beta_i \supsetneq \alpha_{i+1} \neq \alpha_i$. We represent such a path diagrammatically as below

$$\alpha_0 \rightarrowtail \beta_0 \longrightarrow \alpha_1 \rightarrowtail \beta_1 \longrightarrow \cdots \longrightarrow \alpha_k \rightarrowtail \beta_k \longrightarrow \alpha_{k+1}$$

(in the diagram above, $\beta^{(d+1)} \to \alpha^{(d)}$ implies $\beta \supsetneq \alpha$). We say such a path is a *nontrivial closed path* if $k \geq 0$ and $\alpha_{k+1} = \alpha_0$.

Definition 2.3 (Gradient vector field). A gradient vector field on a simplicial complex $\mathcal{K}$ is a discrete vector field $\mathcal{V}$ on $\mathcal{K}$ which does not admit nontrivial closed $\mathcal{V}$-paths.

For a gradient vector field $\mathcal{V}$, when it is clear from the context, we sometimes call a $\mathcal{V}$-path a *gradient path*.

Let $\mathcal{V}$ be a gradient vector field on a simplicial complex $\mathcal{K}$. We call a nonempty simplex α a *critical simplex* (with respect to $\mathcal{V}$) if one of the following holds:

(i) α does not appear in any pair of $\mathcal{V}$, or

(ii) α is a 0-simplex and $(\emptyset, \alpha) \in \mathcal{V}$.

We recall that a *CW complex* is a topological space built recursively by gluing cells (which are homeomorphic copies of balls) of increasing dimension. The fundamental theorem of discrete Morse theory is as below.

Theorem 2.4 (Forman (1998, 2002)). *If $\mathcal{K}$ is a simplicial complex and $\mathcal{V}$ is a gradient vector field on $\mathcal{K}$, then $\mathcal{K}$ is homotopy equivalent to a CW complex with exactly one cell of dimension p for each critical simplex (with respect to $\mathcal{V}$) of dimension p.*

The following is an important corollary of Theorem 2.4.

Theorem 2.5 (Forman (2002)). *If $\mathcal{K}$ is a simplicial complex and $\mathcal{V}$ is a gradient vector field on $\mathcal{K}$ such that the only critical simplices are one 0-simplex and k simplices of dimension d, then $\mathcal{K}$ is homotopy equivalent to the wedge of k spheres of dimension d.*

Theorem 2.4 implies that the topological information pertaining to $\mathcal{K}$ would be concise and easier to compute if the number of critical simplices of each dimension, with respect to $\mathcal{V}$, is as low as possible. If m_i is the number of i-dimensional critical simplices of $\mathcal{K}$, and b_i is the i-th Betti number of $\mathcal{K}$, then we have the following inequalities.

Theorem 2.6 (Morse inequalities, Forman (1998, 2002)).

The weak Morse inequalities: *If d is the dimension of $\mathcal{K}$, then*

(i) for each $i \in \{0, 1, \ldots, d\}$, $m_i \geq b_i$,

(ii) $m_0 - m_1 + \cdots + (-1)^d m_d = b_0 - b_1 + \cdots + (-1)^d b_d$.

The strong Morse inequalities: *For each $i \geq 0$,*

$$m_i - m_{i-1} + \cdots + (-1)^i m_0 \geq b_i - b_{i-1} + \cdots + (-1)^i b_0.$$

We call a gradient vector field a *perfect gradient vector field* if $m_i = b_i$ for all i. Since b_i is the free rank of the i-th homology group of $\mathcal{K}$, it follows that no perfect gradient vector field exists on $\mathcal{K}$ if a homology group of $\mathcal{K}$ has torsion. Moreover, a perfect gradient vector field on a complex may not exist even when all the homology groups are torsion-free, e.g., the dunce hat (Ayala et al. (2012); Whitehead (1939); Zeeman (1963)). This motivates us to call a gradient vector field an *optimal gradient vector field* if the number of critical simplices is the least possible (in comparison with all other gradient vector fields on the same complex). Here we note that the problem of finding an optimal gradient vector field (equivalently, finding a gradient vector field of the highest cardinality) on a given complex is not a computationally easy problem in general; in fact it was shown to be an NP-hard problem (Joswig and Pfetsch (2006); Lewiner et al. (2003a)), even for 2-dimensional complexes (Eğecioğlu and Gonzalez (1996); Lewiner et al. (2003b)). A linear algorithm to find optimal gradient vector fields on (discrete) 2-manifolds was provided by Lewiner et al. (2003b). Although this algorithm can be extended to CW complexes of dimension (up to) two without the manifold property (an example of such a complex, in relation to this article, is the matching complex M_7), the resulting gradient vector field may be arbitrarily far from the optimum.

The following is a useful result to augment a given gradient vector field on a complex (i.e., to reduce the number of critical simplices) by "cancelling" a pair of critical simplices.

Theorem 2.7 (Cancellation of a critical pair, Forman (1998, 2002))**.** *Suppose $\mathcal{V}$ is a gradient vector field on a complex $\mathcal{K}$ such that $\alpha^{(d)}$ and $\beta^{(d+1)}$ are critical. If there is a unique $\mathcal{V}$-path from a d-simplex contained in β, say $\alpha_0^{(d)}$, to α, then there is a gradient vector field $\mathcal{V}'$ on $\mathcal{K}$ such that the critical simplices with respect to $\mathcal{V}'$ remain the same, except that α and β are no longer critical. Moreover, $\mathcal{V}'$ is same as $\mathcal{V}$ except along the unique $\mathcal{V}$-path from α_0 to α.*

A sketch of a proof is as follows. Let the unique $\mathcal{V}$-path from α_0 to α be

$$\alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \ldots, \alpha_k^{(d)}, \beta_k^{(d+1)}, \alpha_{k+1}^{(d)} = \alpha.$$

Thus, we have the following diagram.

$$\beta \longrightarrow \alpha_0 \rightarrowtail \beta_0 \longrightarrow \alpha_1 \rightarrowtail \beta_1 \longrightarrow \cdots \longrightarrow \alpha_k \rightarrowtail \beta_k \longrightarrow \alpha_{k+1} = \alpha$$

We get $\mathcal{V}'$ from $\mathcal{V}$ by reversing the arrows (with $\to$ becoming $\leftarrowtail$ and $\rightarrowtail$ becoming $\leftarrow$) in the diagram above.

$$\beta \leftarrowtail \alpha_0 \longleftarrow \beta_0 \leftarrowtail \alpha_1 \longleftarrow \beta_1 \leftarrowtail \cdots \leftarrowtail \alpha_k \longleftarrow \beta_k \leftarrowtail \alpha_{k+1} = \alpha$$

In other words,

$$\mathcal{V}' = (\mathcal{V} \setminus \{(\alpha_i, \beta_i) : i \in \{0, 1, \ldots, k\}\}) \sqcup \{(\alpha_{i+1}, \beta_i) : i \in \{0, 1, \ldots, k\}\} \sqcup \{(\alpha_0, \beta)\}$$

($\sqcup$ denotes the union of disjoint sets). The uniqueness of the $\mathcal{V}$-path from α_0 to α guarantees that $\mathcal{V}'$ is also a gradient vector field on $\mathcal{K}$. Moreover, it implies that α and β are not critical with respect to $\mathcal{V}'$, while the criticality of all other simplices remains unchanged.

The following allows us to apply the technique above to cancel several pairs of critical simplices simultaneously.

Theorem 2.8 (Hersh (2005))**.** *Let $\mathcal{V}$ be a gradient vector field on a complex $\mathcal{K}$ such that for $i \in \{1, \ldots, r\}$ there is a unique $\mathcal{V}$-path γ_i from a $(d_i - 1)$-simplex contained in the critical d_i-simplex β_i to the critical $(d_i - 1)$-simplex α_i. If there is no non-identity permutation π of r elements such that there is a $\mathcal{V}$-path from a $(d_i - 1)$-simplex contained in β_i to $\alpha_{\pi(i)}$ for all $i \in \{1, \ldots, r\}$, then reversing all the $\mathcal{V}$-paths γ_i (to cancel the critical pair α_i and β_i) would still produce a gradient vector field on $\mathcal{K}$.*

2.4 Discrete Morse homology

Let $\mathcal{K}$ be an abstract simplicial complex and $\mathcal{V}$ be a given gradient vector field on $\mathcal{K}$. Let us fix an ordering on $V(\mathcal{K})$, which induces an orientation on the simplices.

First, we need to introduce the notion of the *incidence number* between two oriented simplices of consecutive dimensions. Let $\beta = [x_0, x_1, \ldots, x_k]$ be a k-simplex. If $\alpha = [x_0, \ldots, \widehat{x_i}, \ldots, x_k]$ is a $(k-1)$-simplex (contained in β), then the incidence number between β and α is $(-1)^i$, and we denote it by $\langle \beta, \alpha \rangle$. Otherwise, if α is a $(k-1)$-simplex such that $\alpha \nsubseteq \beta$, then we define $\langle \beta, \alpha \rangle$ to be 0.

Now, let

$$\gamma : \alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \ldots, \alpha_{k-1}^{(d)}, \beta_{k-1}^{(d+1)}, \alpha_k^{(d)}$$

be a $\mathcal{V}$-path. The *multiplicity* of γ (Forman (1998); Gallais (2010)), denoted by $m(\gamma)$, is given by

$$m(\gamma) := \prod_{i=0}^{k-1} (-1) \langle \beta_i, \alpha_i \rangle \langle \beta_i, \alpha_{i+1} \rangle = (-1)^k \prod_{i=0}^{k-1} \langle \beta_i, \alpha_i \rangle \langle \beta_i, \alpha_{i+1} \rangle. \tag{2.2}$$

We observe that, for all i, both $\langle \beta_i, \alpha_i \rangle$, $\langle \beta_i, \alpha_{i+1} \rangle \in \{-1, +1\}$, and thus $m(\gamma) \in \{-1, +1\}$.

Let $\Gamma(\sigma', \sigma)$ be the set of $\mathcal{V}$-paths starting at σ' and ending at σ. Let us denote the free abelian group generated by all critical (with respect to $\mathcal{V}$) k-simplices of $\mathcal{K}$ by $\tilde{C}_k(\mathcal{K})$. For a (oriented) critical k-simplex τ, first we define a boundary operator $\tilde{\partial}_k$ on τ as below, and then extend it linearly to $\tilde{\partial}_k : \tilde{C}_k(\mathcal{K}) \to \tilde{C}_{k-1}(\mathcal{K})$.

$$\tilde{\partial}_k(\tau) := \sum_{\sigma:\, \text{critical}} n(\tau^{(k)}, \sigma^{(k-1)}) \cdot \sigma^{(k-1)}, \tag{2.3}$$

where

$$n(\tau, \sigma) := \sum_{\sigma' \subsetneq \tau} \langle \tau^{(k)}, \sigma'^{(k-1)} \rangle \sum_{\gamma \in \Gamma(\sigma', \sigma)} m(\gamma). \tag{2.4}$$

It follows that, for all $k \geq 0$, we have $\tilde{\partial}_k \circ \tilde{\partial}_{k+1} = 0$ (Forman (1998); Gallais (2010)). Thus, $(\tilde{C}_*(\mathcal{K}), \tilde{\partial}_*)$ is a chain complex, and we call the homology of this chain complex the (discrete) *Morse homology* of $\mathcal{K}$ (with respect to the chosen gradient vector field $\mathcal{V}$).

Theorem 2.9 (Forman (1998)). *With respect to any given gradient vector field $\mathcal{V}$ on $\mathcal{K}$, the chain complex $(\tilde{C}_*(\mathcal{K}), \tilde{\partial}_*)$ is homotopy equivalent to the simplicial chain complex $(C_*(\mathcal{K}), \partial_*)$. Consequently, the discrete Morse homology, which is independent of the chosen gradient vector field, is isomorphic to the simplicial homology, i.e.,*

$$H_k(\tilde{C}_*(\mathcal{K}), \tilde{\partial}_*) = H_k(\mathcal{K}), \text{ for all } k \geq 0.$$

3   Construction of a gradient vector field on M_n

In this section, we construct a gradient vector field on M_n for $n \geq 5$, which doesn't admit any critical simplices of dimension up to (and including) $\nu_n - 1$, except one unavoidable critical 0-simplex. For the rest of the article, we implicitly assume $n \geq 5$ whenever we talk about M_n in general.

For a positive integer k, we denote the set $\{1, \ldots, k\}$ by $[k]$. We partition the vertex set $V(K_n)$ into $\lceil \frac{n}{3} \rceil$ sets, and label the vertices of K_n depending on the part they belong to as below.

1. For $n = 3m$, $V(K_n) = V_1 \sqcup \ldots \sqcup V_m$, where $V_i = \{v_1^{(i)}, v_2^{(i)}, v_3^{(i)}\}$ for all $i \in [m]$.

2. For $n = 3m + 1$, $V(K_n) = V_1 \sqcup \ldots \sqcup V_{m+1}$, where $V_i = \{v_1^{(i)}, v_2^{(i)}, v_3^{(i)}\}$ for all $i \in [m]$, and $V_{m+1} = \{v_1^{(m+1)}\}$.

3. For $n = 3m + 2$, $V(K_n) = V_1 \sqcup \ldots \sqcup V_{m+1}$, where $V_i = \{v_1^{(i)}, v_2^{(i)}, v_3^{(i)}\}$ for all $i \in [m]$, and $V_{m+1} = \{v_1^{(m+1)}, v_2^{(m+1)}\}$.

We call an edge e an *i-level edge* (or simply, a *levelled-edge*) if both of its endvertices are in the same V_i for some i. Otherwise, if endvertices of e are in V_i and V_j with $i \neq j$, then e is called a *cross-edge* (between V_i and V_j) (see Fig. 2).

First, we define the following discrete vector fields on M_n:

$$\mathcal{M}_1' = \left\{ \left(\alpha, \alpha \sqcup \{v_j^{(1)} v_k^{(1)}\} \right) : \alpha \in M_n, \alpha \text{ covers only } v_i^{(1)} \text{ of } V_1, \{i, j, k\} = [3] \right\},$$

$$\mathcal{M}_1'' = \left\{ \left(\alpha, \alpha \sqcup \{v_2^{(1)} v_3^{(1)}\} \right) : \alpha \in M_n, \alpha \text{ leaves entire } V_1 \text{ uncovered} \right\}$$
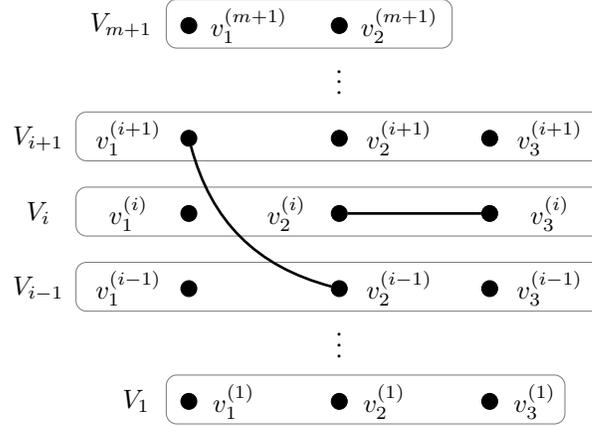
Fig. 2: The partition of $V(K_{3m+2})$ into $m + 1$ levels. Here, $v_2^{(i)} v_3^{(i)}$ is an i-level edge and $v_2^{(i-1)} v_1^{(i+1)}$ is a cross-edge between V_{i-1} and V_{i+1}.
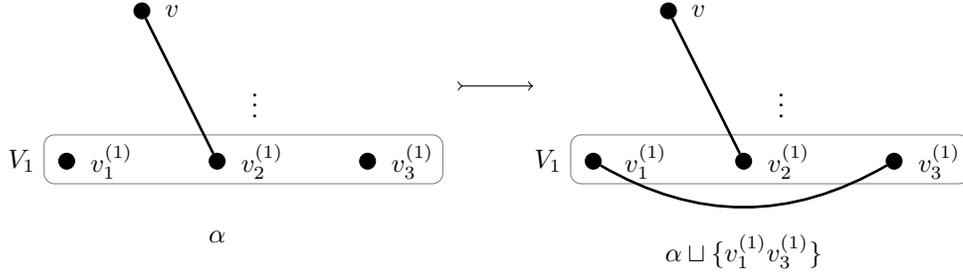
(see Fig. 3 and Fig. 4).

Fig. 3: The matching α covers only $v_2^{(1)}$ of V_1. Thus, α is paired off with $\alpha \sqcup \{v_1^{(1)} v_3^{(1)}\}$ in $\mathcal{M}_1'$.
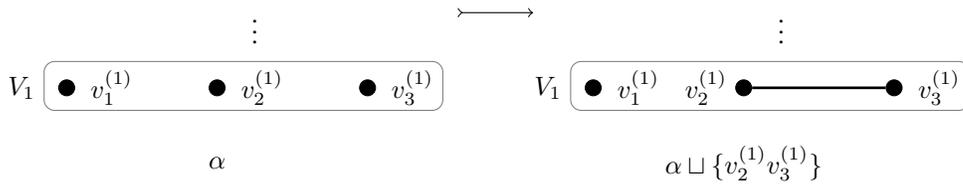

Fig. 4: The matching α leaves entire V_1 uncovered. Thus, α is paired off with $\alpha \sqcup \{v_2^{(1)} v_3^{(1)}\}$ in $\mathcal{M}_1''$.

Let $\mathcal{M}_1 = \mathcal{M}_1' \sqcup \mathcal{M}_1''$. We note that $\mathcal{M}_1$ is also a discrete vector field on M_n. Let $\mathcal{U}_1$ denote the set of all matchings in K_n that are not paired off in $\mathcal{M}_1$, i.e., $\mathcal{U}_1 = \{\alpha \in M_n : \alpha$ does not appear in any pair of $\mathcal{M}_1\}$. We observe that $\alpha \in \mathcal{U}_1$ if and only if one of the following holds:

(i) at least two vertices of V_1 are matched by α with vertices outside V_1,

(ii) α contains the edge $v_1^{(1)}v_i^{(1)}$, where $i \in \{2,3\}$ and α doesn't cover the vertex $v \in V_1 \setminus \{v_1^{(1)}, v_i^{(1)}\}$
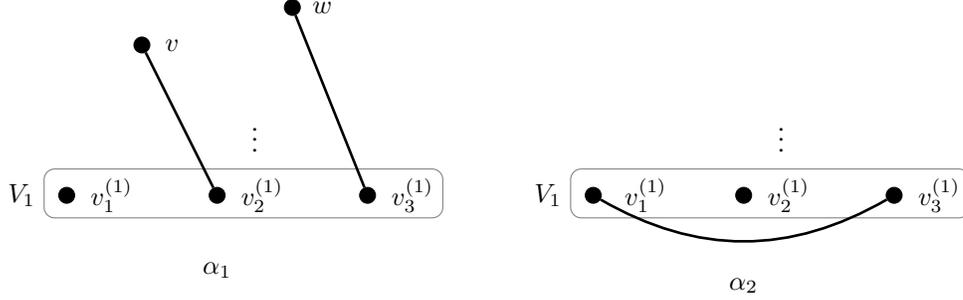
(see Fig. 5).

Fig. 5: The matching α_1 matches $v_2^{(1)}$ and $v_3^{(1)}$ of V_1 with vertices outside V_1, whereas α_2 contains the edge $v_1^{(1)}v_3^{(1)}$, but leaves $v_2^{(1)}$ uncovered. Thus, both $\alpha_1, \alpha_2 \in \mathcal{U}_1$.

Next, following the same scheme as before, we define the following discrete vector fields on M_n:

$$\mathcal{M}_2' = \left\{ \left(\alpha, \alpha \sqcup \{v_j^{(2)}v_k^{(2)}\}\right) : \alpha \in \mathcal{U}_1, \alpha \text{ covers only } v_i^{(2)} \text{ of } V_2, \{i,j,k\} = [3] \right\},$$

$$\mathcal{M}_2'' = \left\{ \left(\alpha, \alpha \sqcup \{v_2^{(2)}v_3^{(2)}\}\right) : \alpha \in \mathcal{U}_1, \alpha \text{ leaves entire } V_2 \text{ uncovered} \right\}.$$

Let $\mathcal{M}_2 = \mathcal{M}_2' \sqcup \mathcal{M}_2''$ (which is also a discrete vector field on M_n) and $\mathcal{U}_2 = \{\alpha \in M_n : \alpha \text{ does not appear in any pair of } \mathcal{M}_1 \sqcup \mathcal{M}_2\}$. Suppose $\alpha \in \mathcal{U}_1$. Then $\alpha \in \mathcal{U}_2$ if and only if one of the following holds:

(i) at least two vertices of V_2 are matched by α with vertices outside V_2,

(ii) α contains the edge $v_1^{(2)}v_i^{(2)}$, where $i \in \{2,3\}$ and α doesn't cover the vertex $v \in V_2 \setminus \{v_1^{(2)}, v_i^{(2)}\}$.

In general, following the same scheme as above, we get a sequence of discrete vector fields $\mathcal{M}_1, \ldots, \mathcal{M}_m$ on M_n and a family of subsets $\mathcal{U}_1 \supseteq \mathcal{U}_2 \supseteq \ldots \supseteq \mathcal{U}_m$ of M_n such that $\mathcal{U}_k = \{\alpha \in M_n : \alpha \text{ does not appear in any pair of } \mathcal{M}_1 \sqcup \ldots \sqcup \mathcal{M}_k\}$, for all $k \in [m]$. We note that for any $k \in \{2, \ldots, m\}$, the discrete vector field $\mathcal{M}_k$ may contain only pairs of matchings of the form $(\alpha, \alpha \sqcup \{e\})$ where both α and $\alpha \sqcup \{e\}$ are in $\mathcal{U}_{k-1}$.

If $n = 3m + 2$, we define another discrete vector field on M_n, viz.,

$$\mathcal{M}_{m+1} = \left\{ \left(\alpha, \alpha \sqcup \{v_1^{(m+1)}v_2^{(m+1)}\}\right) : \alpha \in \mathcal{U}_m, \alpha \text{ covers neither } v_1^{(m+1)} \text{ nor } v_2^{(m+1)} \right\}.$$

For $n = 3m$ or $3m + 1$, let $\mathcal{M} = \mathcal{M}_1 \sqcup \ldots \sqcup \mathcal{M}_m$ and for $n = 3m + 2$, let $\mathcal{M} = \mathcal{M}_1 \sqcup \ldots \sqcup \mathcal{M}_{m+1}$. We note that $\mathcal{M}$ as defined above is a collection of pairs of matchings in K_n, and thus $\mathcal{M}$ also depends on n. To make the notation less cumbersome, we avoid adding the parameter n to $\mathcal{M}$.

From the construction of $\mathcal{M}$, we make the following observation.

Observation 3.1. *For any n, the following hold.*

1. *$\mathcal{M}$ is a discrete vector field on M_n.*

2. *If $(\alpha, \alpha \sqcup \{e\}) \in \mathcal{M}$, then e is an i-level edge, for some $i \in [\lceil \frac{n}{3} \rceil]$.*

We now prove that $\mathcal{M}$ is a gradient vector field on M_n.

Proposition 3.2. *For all n, the discrete vector field $\mathcal{M}$ is a gradient vector field on M_n.*

Proof: Let, if possible, $\alpha_0, \beta_0, \alpha_1, \beta_1, \ldots, \alpha_r, \beta_r, \alpha_{r+1} = \alpha_0$ be a nontrivial closed $\mathcal{M}$-path, and let us denote it by γ. Let $\beta_0 = \alpha_0 \sqcup \{e_0\}$ and $\alpha_1 = \beta_0 \setminus \{e_0'\}$ (with $e_0 \neq e_0'$).

$$\alpha_0 \xrightarrow{\sqcup \{e_0\}} \beta_0 \xrightarrow{-\{e_0'\}} \alpha_1 \longmapsto \beta_1 \longrightarrow \cdots \longrightarrow \alpha_r \longmapsto \beta_r \longrightarrow \alpha_{r+1} = \alpha_0$$

From Observation 3.1, it follows that e_0 is a levelled-edge. Now, if e_0' is a cross-edge, then from the construction of $\mathcal{M}$, it follows that $e_0' \notin \alpha_p$, for all $p \geq 1$. This contradicts the assumption that γ is a nontrivial closed $\mathcal{M}$-path as $e_0' \in \alpha_0$. Thus, e_0' is also a levelled-edge.

Let e_0 and e_0' be an i_0-level edge and an i_1-level edge for some $i_0, i_1 \in [\lceil \frac{n}{3} \rceil]$, respectively. Thus, it follows that $(\alpha_0, \beta_0) \in \mathcal{M}_{i_0}$ and both $\alpha_0, \beta_0 \in \mathcal{U}_k$ for all $k < i_0$.

We note that $i_0 \neq i_1$ as both $e_0, e_0' \in \beta_0$. Now, if $i_1 > i_0$, then we observe that α_1 matches the vertices of $V_1 \sqcup \ldots \sqcup V_{i_0}$ in *exactly the same manner* as β_0 does. This implies, just as β_0, the matching $\alpha_1 \in \mathcal{U}_k$ for all $k < i_0$. Since $(\alpha_0 = \beta_0 \setminus \{e_0\}, \beta_0) \in \mathcal{M}$, we have $(\alpha_1 \setminus \{e_0\}, \alpha_1) \in \mathcal{M}$, a contradiction as $(\alpha_1, \beta_1) \in \mathcal{M}$. Therefore, we have $i_1 < i_0$.

If $e_0' = v_k^{(i_1)} v_\ell^{(i_1)}$ (where $k, \ell \in [3]$) and the vertex $v_m^{(i_1)}$ (where $m \in [3] \setminus \{k, \ell\}$) is matched by α_0 (with a vertex outside V_{i_1}), then from the properties of simplices appearing in $\mathcal{U}_{i_1}$, we have $\alpha_0 \notin \mathcal{U}_{i_1}$, a contradiction. Similarly, if $e_0' = v_2^{(i_1)} v_3^{(i_1)}$ and the vertex $v_1^{(i_1)}$ is not covered α_0, then again we have $\alpha_0 \notin \mathcal{U}_{i_1}$, a contradiction.

So suppose $e_0' = v_1^{(i_1)} v_k^{(i_1)}$ (where $k \in \{2, 3\}$) and the vertex $v_\ell^{(i_1)}$ (where $\ell \in [3] \setminus \{1, k\}$) is not covered by α_0, and thus not covered by α_1 as well. In this case, $\beta_1 = \alpha_1 \sqcup \{v_2^{(i_1)} v_3^{(i_1)}\}$ and $(\alpha_1, \beta_1) \in \mathcal{M}_{i_1}$ (we note that β_1 doesn't contain e_0', and thus neither does α_2). By a similar argument as before, α_2 is of the form $\beta_1 \setminus \{e_1'\}$ and β_2 is of the form $\alpha_2 \sqcup \{e_2\}$, where e_1' and e_2 are two i_2-level edges (distinct from each other) with $i_2 < i_1$. Thus, $e_2 \neq e_0'$, and consequently $e_0' \notin \beta_2, \alpha_3$.

$$\alpha_0 \xrightarrow[\;(i_0)\;]{\sqcup \{e_0\}} \beta_0 \xrightarrow[\;(i_1)\;]{-\{e_0' = v_1^{(i_1)} v_k^{(i_1)}\}} \alpha_1 \xrightarrow[\;(i_1)\;]{\sqcup \{v_2^{(i_1)} v_3^{(i_1)}\}} \beta_1 \xrightarrow[\;(i_2)\;]{-\{e_1'\}} \alpha_2 \xrightarrow[\;(i_2)\;]{\sqcup \{e_2\}} \beta_2 \longrightarrow \cdots \longrightarrow \alpha_{r+1}$$
$$\text{(with } i_0 > i_1 > i_2 > \cdots \text{ and so on)}$$

By an inductive argument, we conclude that $e_0' \notin \alpha_p, \beta_p$, for all $p \geq 1$. Thus, $\alpha_{r+1} \neq \alpha_0$ as $e_0' \in \alpha_0$, which is a contradiction. $\qquad\blacksquare$

We now show that the constructed gradient vector field satisfies the requirements of Theorem 1.2.

Proof of Theorem 1.2: We claim that the gradient vector field $\mathcal{M}$ on M_n doesn't admit critical simplices of dimension up to $\nu_n - 1$ for any n, except one 0-simplex. From the construction of $\mathcal{M}$ on M_n, we observe that if $\alpha \in M_n$ is not paired off in $\mathcal{M}$, then for each $i \in \lceil \frac{n}{3} \rceil$, the matching α leaves at most one vertex of V_i uncovered. Therefore, α leaves at most $\lceil \frac{n}{3} \rceil$ vertices of the graph K_n uncovered. This implies

$$|\alpha| \geq \frac{1}{2} \left(n - \left\lceil \frac{n}{3} \right\rceil \right) > \left\lfloor \frac{n+1}{3} \right\rfloor - 1 = \nu_n,$$

i.e., the dimension of α is strictly greater than $\nu_n - 1$.

Also, the 0-simplex $\xi := \{v_2^{(1)}v_3^{(1)}\}$ is paired off with $\emptyset$ in $\mathcal{M}$, and thus ξ is the only critical 0-simplex. Therefore, all simplices of M_n of dimension up to $\nu_n - 1$, except ξ, are not critical with respect to $\mathcal{M}$. $\square$

By Theorem 2.4, it follows that M_n is homotopy equivalent to a CW complex with no cells of dimension up to $\nu_n - 1$, except one 0-cell. Thus, Theorem 1.1 follows as a corollary of Theorem 1.2. However, we remark that the existence of such a gradient vector field is in fact a stronger notion. Even when a complex is homotopically k-connected, it is not guaranteed that there is a gradient vector field on it, with no critical simplices (except one 0-simplex) of dimension up to k. A well-known example is the *dunce hat*. It is a contractible space (and thus simply connected), but no gradient vector field, with a 0-simplex as the only critical simplex, can be assigned on any triangulation of the dunce hat (Ayala et al. (2012); Whitehead (1939); Zeeman (1963)).

Example 3.3 (Homotopy type of M_8). In order to determine the homotopy type of M_8, we consider the gradient vector field $\mathcal{M}$ defined on M_8, and extend it to a perfect gradient vector field as follows. Let α° be a 3-simplex in M_8 (i.e., α° is a perfect matching in K_8), which is critical with respect to $\mathcal{M}$. From the construction of $\mathcal{M}$, it follows that all the edges in α° are cross-edges. Moreover, for any $e \in \alpha^\circ$, the 2-simplex $\alpha^\circ \setminus \{e\}$ is also critical with respect to $\mathcal{M}$. Now, since α° is a perfect matching containing only cross-edges, α° matches $v_1^{(3)}$ with a vertex outside V_3, say v°. So v° is of the form $v_j^{(i)}$, for some $i \in [2]$ and $j \in [3]$. We extend $\mathcal{M}$ to a discrete vector field $\mathcal{M}^\circ$ by adding the new pairs of the form $(\alpha^\circ \setminus \{v_1^{(3)}v^\circ\}, \alpha^\circ)$ to $\mathcal{M}$ (see Fig. 6), for each 3-simplex α° which is critical with respect to $\mathcal{M}$.
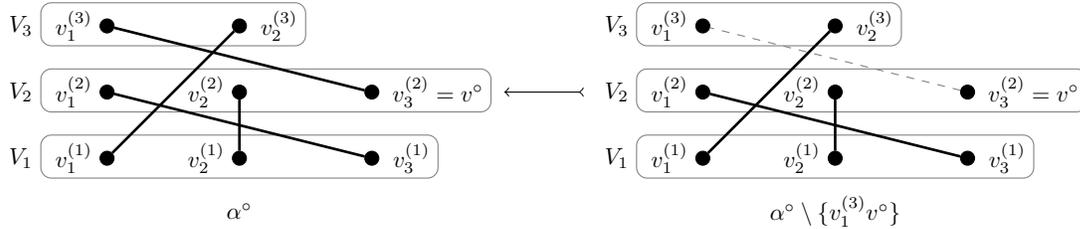

Fig. 6: The critical (with respect to $\mathcal{M}$) 3-simplex α° is paired off with the critical 2-simplex $\alpha^\circ \setminus \{v_1^{(3)}v^\circ\}$ in $\mathcal{M}^\circ$.

Any $\mathcal{M}^\circ$-path γ, which is not an $\mathcal{M}$-path, contains a pair $(\alpha_i^{(2)}, \beta_i^{(3)}) \in \mathcal{M}^\circ \setminus \mathcal{M}$. If $\alpha_{i+1} (\neq \alpha_i)$ is any 2-simplex contained in β_i, then γ does not extend beyond such an α_{i+1}, as β_i is the only 3-simplex containing α_{i+1}. This implies $\mathcal{M}^\circ$ is also a gradient vector field on M_8, and moreover, all the 0-simplices, 1-simplices, and 3-simplices of M_8 are paired off in $\mathcal{M}^\circ$. Thus, only critical simplices of M_8 (with respect to $\mathcal{M}^\circ$) are some 2-simplices and exactly one 0-simplex, viz., the matching $\{v_2^{(1)}v_3^{(1)}\}$. The number of critical 2-simplices may be determined from the Euler characteristic of M_8 (or, alternatively, by direct counting). If $(f_0, \ldots, f_3)$ is the f-vector of M_8, then from Equation (2.1), $f_0 = 28$, $f_1 = 210$, $f_2 = 420$, $f_3 = 105$, and thus $\chi(M_8) = f_0 - f_1 + f_2 - f_3 = 133$. Therefore, by Theorem 2.5, we conclude that M_8 is homotopy equivalent to a wedge of 132 spheres of dimension 2.

4 Homology groups of M_7

4.1 Construction of a near-optimal gradient vector field on M_7

We now consider the matching complex M_7 as an example, and determine its homotopy type using techniques developed in this article. Throughout the rest of this article, $\mathcal{M}$ stands for the gradient vector field on M_7 in particular, as constructed in Section 3.

We extend the gradient vector field $\mathcal{M}$ to a more useful one as follows. Let α^* be a critical (with respect to $\mathcal{M}$) 1-simplex containing two cross-edges between V_1 and V_2. Let v^* be only vertex of V_1 that is left uncovered by α^*. We observe that $\alpha^* \sqcup \{v_1^{(3)} v^*\}$ is a critical 2-simplex with respect to $\mathcal{M}$. Thus, we may extend the gradient vector field $\mathcal{M}$ to the discrete vector field $\mathcal{M}^*$ by adding the new pairs of the form $(\alpha^*, \alpha^* \sqcup \{v_1^{(3)} v^*\})$ to $\mathcal{M}$ (see Fig. 7), for each critical (with respect to $\mathcal{M}$) 1-simplex α^* containing two cross-edges between V_1 and V_2.
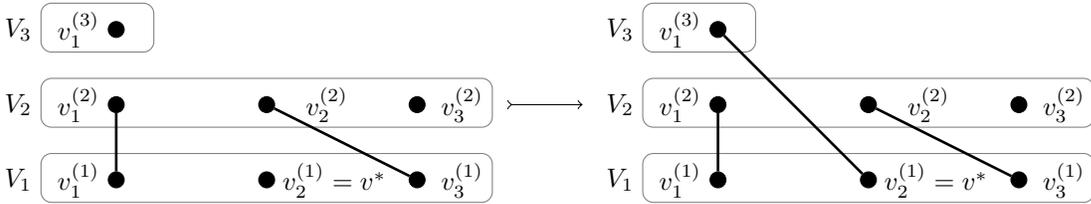
Fig. 7: The critical (with respect to $\mathcal{M}$) 1-simplex $\alpha^* = \{v_1^{(1)} v_1^{(2)}, v_3^{(1)} v_2^{(2)}\}$ is paired off with the critical 2-simplex $\alpha^* \sqcup \{v_1^{(3)} v_2^{(1)}\}$ in $\mathcal{M}^*$.

Proposition 4.1. *The discrete vector field $\mathcal{M}^*$ is a gradient vector field on M_7.*

Proof: Let, if possible, $\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \ldots, \alpha_k^{(p)}, \beta_k^{(p+1)}, \alpha_{k+1}^{(p)} = \alpha_0^{(p)}$ be a nontrivial closed $\mathcal{M}^*$-path. If $p = 0$, then it leads to a contradiction to the fact that $\mathcal{M}$ is a gradient vector field. So let $p = 1$.

If $(\alpha_i, \beta_i) \in \mathcal{M}$ for all $i \in \{0, 1, \ldots, k\}$, then again it contradicts the fact that $\mathcal{M}$ is a gradient vector field. So without loss of generality, let $(\alpha_0, \beta_0) \in \mathcal{M}^* \setminus \mathcal{M}$. In this case, α_0 consists of two cross-edges between V_1 and V_2. Let $\alpha_1 = \beta_0 \setminus \{e'\}$, where e' is an edge in α_0. We note that if $\beta = \alpha \sqcup \{e\}$ and $(\alpha, \beta) \in \mathcal{M}^*$, then e is either a levelled-edge or a cross-edge between V_1 and V_3. This implies, for all $i \in [k]$, $e' \notin \beta_i$, and subsequently $e' \notin \alpha_{i+1}$. In particular, $e' \notin \alpha_{k+1}$. This leads to a contradiction, as $\alpha_{k+1} = \alpha_0$. $\qquad\square$

4.2 Computation of (Morse) homology groups of M_7

Throughout this subsection, we always assume the gradient vector field $\mathcal{M}^*$ on the matching complex M_7 while discussing the nature of a simplex (i.e., criticality and other related notions). Also, hereafter, while representing a simplex of M_7 (i.e., a matching in K_7) diagrammatically, we would not explicitly label the vertices of K_7 and the sets V_1, V_2, and V_3 for the sake of simplicity, and so it should be understood from the context. For example, we would represent the matching $\{v_1^{(1)} v_1^{(2)}, v_2^{(1)} v_1^{(3)}, v_3^{(1)} v_2^{(2)}\}$ (see Fig. 7) by

the following.

Observation 4.2 (Characterization of critical simplices)**.**

1. *The only critical 0-simplex is $\xi = \{v_2^{(1)} v_3^{(1)}\}$ (thus $\tilde{C}_0(M_7) \cong \mathbb{Z}$).*

2. *Critical 1-simplices are $\sigma_1 := \{v_1^{(1)} v_2^{(1)}, v_1^{(2)} v_2^{(2)}\}$, $\sigma_2 := \{v_1^{(1)} v_2^{(1)}, v_1^{(2)} v_3^{(2)}\}$, $\sigma_3 := \{v_1^{(1)} v_3^{(1)}, v_1^{(2)} v_2^{(2)}\}$, and $\sigma_4 := \{v_1^{(1)} v_3^{(1)}, v_1^{(2)} v_3^{(2)}\}$.*

3. *Critical 2-simplices are of the form $\{e_1, e_2, e_3\}$ with one of the following.*

   (a) *Each of e_1, e_2, and e_3 is a cross-edge between V_1 and V_2 (there are 6 such simplices).*

   (b) *Two of e_1, e_2, and e_3 are cross-edges between V_1 and V_2, and the remaining one is a cross-edge between V_2 and V_3 (there are 18 such simplices).*

   Therefore, there are 24 critical 2-simplices.

First, we assign a unique label $\ell(e)$ on the edge $e = v_{j_1}^{(i_1)} v_{j_2}^{(i_2)}$ as follows.

$$\ell(e) = \begin{cases} i_1 j_1 i_2 j_2, & \text{if } i_1 < i_2 \\ i_2 j_2 i_1 j_1, & \text{if } i_2 < i_1 \\ i_1 j_1 i_2 j_2, & \text{if } i_1 = i_2, j_1 < j_2 \\ i_2 j_2 i_1 j_1, & \text{if } i_1 = i_2, j_2 < j_1 \end{cases}$$

Next, we introduce a total order $\leq$ on the vertex set of the complex M_7 (i.e., on $E(K_7)$) by declaring $e_1 \leq e_2$ if and only if $\ell(e_1) \leq \ell(e_2)$ in the *lexicographic order*. We assign each simplex the orientation induced by this total order on $V(M_7)$, i.e., if $\alpha = \{e_0, e_1, \ldots, e_k\}$ is a matching in K_7 with $e_0 < e_1 < \cdots < e_k$, then we denote the oriented k-simplex $[e_0, e_1, \ldots, e_k]$ also by α whenever needed.

4.2.1   Kernel and image of $\tilde{\partial}_1$

If σ is a critical 1-simplex and $\alpha^{(0)} \subsetneq \sigma$, then α is one of $\{v_1^{(1)} v_2^{(1)}\}$, $\{v_1^{(1)} v_3^{(1)}\}$, $\{v_1^{(2)} v_2^{(2)}\}$, and $\{v_1^{(2)} v_3^{(2)}\}$. We note that one of the following two cases holds.

Case 1: $\alpha = \{v_1^{(1)} v_i^{(1)}\}$, where $i \in \{2, 3\}$

The only possible $\mathcal{M}^*$-path that starts from α, and ends at $\xi = \{v_2^{(1)} v_3^{(1)}\}$ is the following.

$$\alpha \longmapsto \{v_1^{(1)} v_i^{(1)}, v_2^{(2)} v_3^{(2)}\} \longrightarrow \{v_2^{(2)} v_3^{(2)}\} \longmapsto \{v_2^{(1)} v_3^{(1)}, v_2^{(2)} v_3^{(2)}\} \longrightarrow \xi$$

Case 2: $\alpha = \{v_1^{(2)} v_i^{(2)}\}$, where $i \in \{2, 3\}$

The only possible $\mathcal{M}^*$-path that starts from α, and ends at $\xi = \{v_2^{(1)} v_3^{(1)}\}$ is the following.

$$\alpha \longmapsto \{v_2^{(1)} v_3^{(1)}, v_1^{(2)} v_i^{(2)}\} \longrightarrow \xi$$

Let us consider the critical 1-simplex $\sigma_1 = \{v_1^{(1)}v_2^{(1)}, v_1^{(2)}v_2^{(2)}\}$. Let γ_1 be the unique $\mathcal{M}^*$-path that starts from $\{v_1^{(1)}v_2^{(1)}\}$ ($\subsetneq \sigma_1$), and ends at $\xi = \{v_2^{(1)}v_3^{(1)}\}$ (from Case 1 above), i.e.,

$$\gamma_1 : \alpha_0 = \{v_1^{(1)}v_2^{(1)}\} \xrightarrow{(-1)} \beta_0 \xrightarrow{(+1)} \alpha_1 \xrightarrow{(+1)} \beta_1 \xrightarrow{(-1)} \alpha_2 = \xi,$$

where $\beta_0 = \{v_1^{(1)}v_2^{(1)}, v_2^{(2)}v_3^{(2)}\}$, $\alpha_1 = \{v_2^{(2)}v_3^{(2)}\}$, $\beta_1 = \{v_2^{(1)}v_3^{(1)}, v_2^{(2)}v_3^{(2)}\}$. Here we have also included $\langle \beta_i, \alpha_j \rangle$, i.e., the incidence number between the oriented simplices β_i and α_j (with $i \leq j \leq i+1$), above the arrow connecting β_i and α_j. Considering α_0, β_0, α_1, β_1, and α_2 as oriented simplices, the multiplicity of γ_1 (from Equation (2.2)),

$$m(\gamma_1) = (-\langle \beta_0, \alpha_0 \rangle \langle \beta_0, \alpha_1 \rangle)(-\langle \beta_1, \alpha_1 \rangle \langle \beta_1, \alpha_2 \rangle)$$
$$= (-(-1)(+1))(-(+1)(-1)) = +1.$$

Let γ_2 be the unique $\mathcal{M}^*$-path that starts from $\{v_1^{(2)}v_2^{(2)}\}$ ($\subsetneq \sigma_1$), and ends at $\xi = \{v_2^{(1)}v_3^{(1)}\}$ (from Case 2 above), i.e.,

$$\gamma_2 : \alpha_0' = \{v_1^{(2)}v_2^{(2)}\} \xrightarrow{(+1)} \beta_0' = \{v_2^{(1)}v_3^{(1)}, v_1^{(2)}v_2^{(2)}\} \xrightarrow{(-1)} \alpha_1' = \xi.$$

Considering α_0', β_0', and α_1' as oriented simplices, the multiplicity of γ_2 (from Equation (2.2)),

$$m(\gamma_2) = -\langle \beta_0', \alpha_0' \rangle \langle \beta_0', \alpha_1' \rangle = -(+1)(-1) = +1.$$

Therefore, from Equation (2.3) and Equation (2.4),

$$\tilde{\partial}_1(\sigma_1) = (\langle \sigma_1, \alpha_0 \rangle \cdot m(\gamma_1) + \langle \sigma_1, \alpha_0' \rangle \cdot m(\gamma_2)) \cdot \xi$$
$$= ((-1)(+1) + (+1)(+1)) \cdot \xi = 0.$$

Analogous computations yield $\tilde{\partial}_1(\sigma_2) = 0$, $\tilde{\partial}_1(\sigma_3) = 0$, and $\tilde{\partial}_1(\sigma_4) = 0$. Thus,

$$\ker(\tilde{\partial}_1) = \left\langle \sigma_1, \sigma_2, \sigma_3, \sigma_4 \right\rangle = \tilde{C}_1(M_7) \text{ and } \operatorname{im}(\tilde{\partial}_1) = 0.$$

Consequently, the zeroth (Morse) homology group of M_7 is $\mathbb{Z}$.

4.2.2   *Image of $\tilde{\partial}_2$*

To begin with, let us consider a critical 2-simplex and compute its image under the boundary operator $\tilde{\partial}_2$ in the following example.

Example 4.3. Let us consider the critical 2-simplex $\eta_1 := \{v_2^{(1)}v_2^{(2)}, v_3^{(1)}v_3^{(2)}, v_1^{(2)}v_1^{(3)}\}$ as shown below.

$$\eta_1 = $$

Let us determine all possible $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_1, and end at a critical 1-simplex. First, a 1-simplex contained in η_1 is one of the following three.

$$\psi_{11} = \qquad \psi_{12} = \qquad \psi_{13} =$$

In Fig. 8, we describe all possible $\mathcal{M}^*$-paths that start from ψ_{11}, and end either at a critical 1-simplex or at a 1-simplex that is paired off with a 0-simplex in $\mathcal{M}^*$.
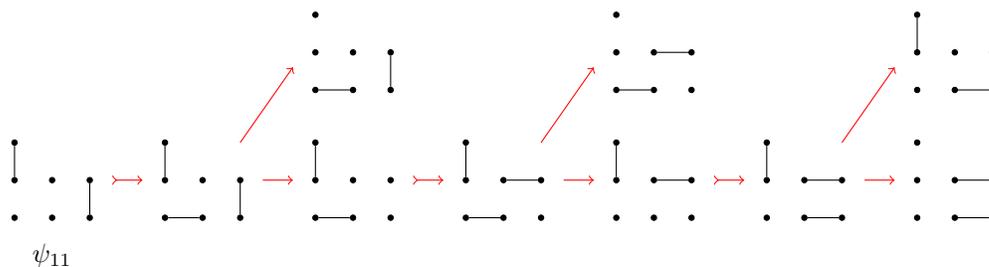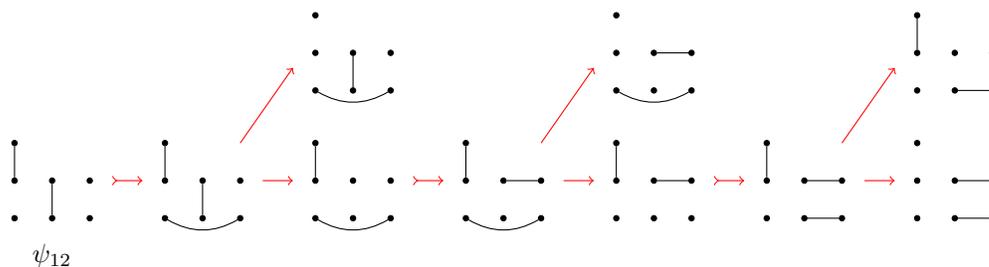

ψ_{11}

Fig. 8: All possible (maximal) $\mathcal{M}^*$-paths that start from ψ_{11}. Note that each of them ends at a 1-simplex that is paired off with a 0-simplex in $\mathcal{M}^*$.

In Fig. 9, we describe all possible $\mathcal{M}^*$-paths that start from ψ_{12}, and end either at a critical 1-simplex or at a 1-simplex that is paired off with a 0-simplex in $\mathcal{M}^*$.

ψ_{12}

Fig. 9: All possible (maximal) $\mathcal{M}^*$-paths that start from ψ_{12}. Note that each of them ends at a 1-simplex that is paired off with a 0-simplex in $\mathcal{M}^*$.

In Fig. 10, we describe all possible $\mathcal{M}^*$-paths that start from ψ_{13}, and end either at a critical 1-simplex or at a 1-simplex that is paired off with a 0-simplex in $\mathcal{M}^*$.

We observe that there are exactly two $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_1, and end at a critical 1-simplex as shown in Fig. 11 (see also Fig. 14 and Fig. 23 in Appendix A).

Let γ_1 be the unique $\mathcal{M}^*$-path that starts from ψ_{13}, and ends at σ_4 (see Fig. 10 and Fig. 11), i.e.,
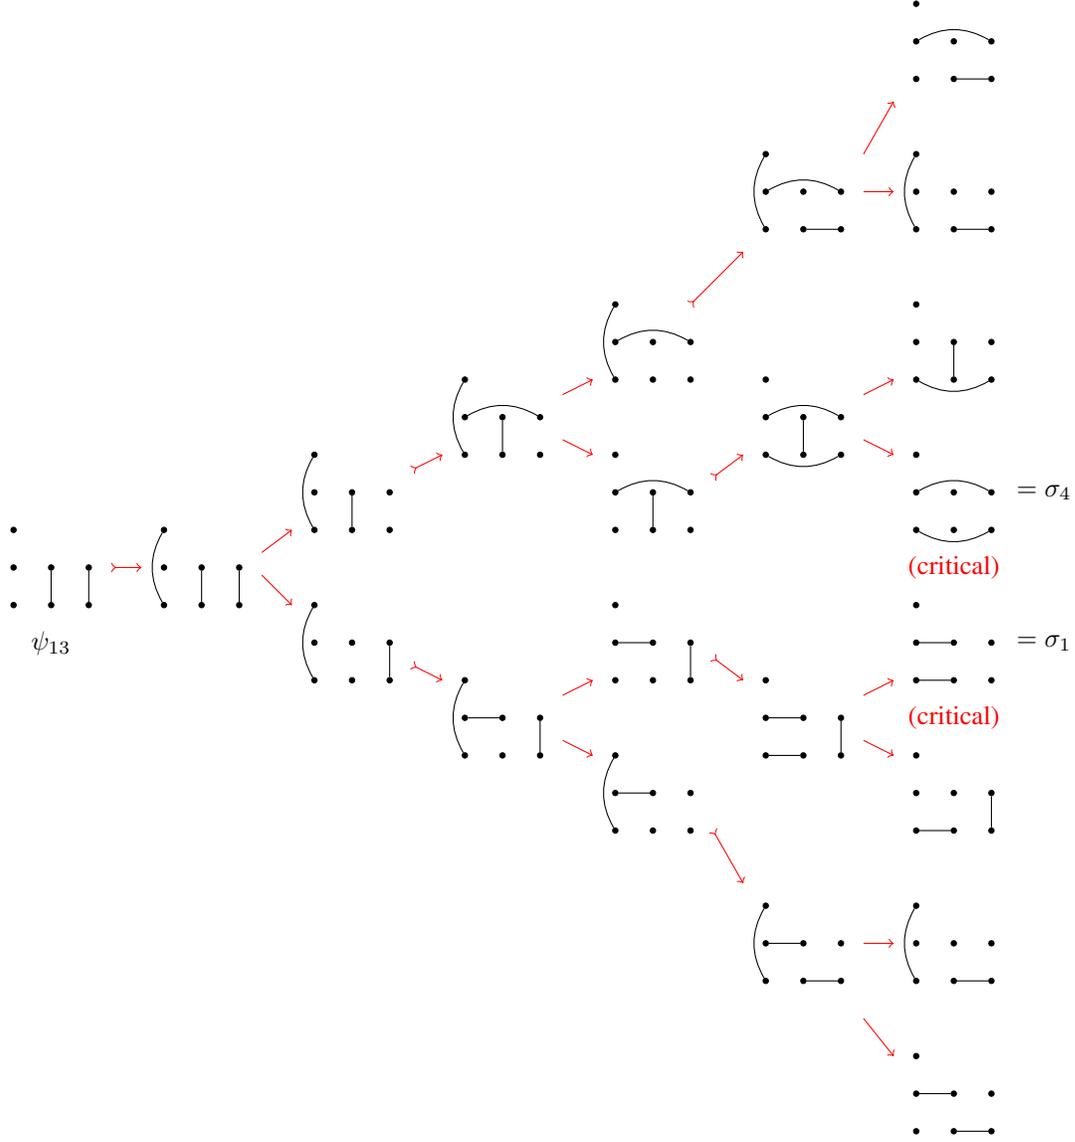
Fig. 10: All possible (maximal) $\mathcal{M}^*$-paths that start from ψ_{13}. Note that only two of them end at a critical 1-simplex.

$$\gamma_1 : \psi_{13} = \alpha_0 \xrightarrow{(+1)} \beta_0 \xrightarrow{(+1)} \alpha_1 \xrightarrow{(+1)} \beta_1 \xrightarrow{(+1)} \alpha_2 \xrightarrow{(+1)} \beta_2 \xrightarrow{(-1)} \alpha_3 = \sigma_4, \text{ where}$$

$$\alpha_0 = [v_2^{(1)} v_2^{(2)}, v_3^{(1)} v_3^{(2)}] = \psi_{13}, \qquad \beta_0 = [v_1^{(1)} v_1^{(3)}, v_2^{(1)} v_2^{(2)}, v_3^{(1)} v_3^{(2)}]$$

$$\alpha_1 = [v_1^{(1)} v_1^{(3)}, v_2^{(1)} v_2^{(2)}], \qquad \beta_1 = [v_1^{(1)} v_1^{(3)}, v_2^{(1)} v_2^{(2)}, v_1^{(2)} v_3^{(2)}]$$

$$\alpha_2 = [v_2^{(1)} v_2^{(2)}, v_1^{(2)} v_3^{(2)}], \qquad \beta_2 = [v_1^{(1)} v_3^{(1)}, v_2^{(1)} v_2^{(2)}, v_1^{(2)} v_3^{(2)}]$$

$$\alpha_3 = [v_1^{(1)} v_3^{(1)}, v_1^{(2)} v_3^{(2)}] = \sigma_4.$$
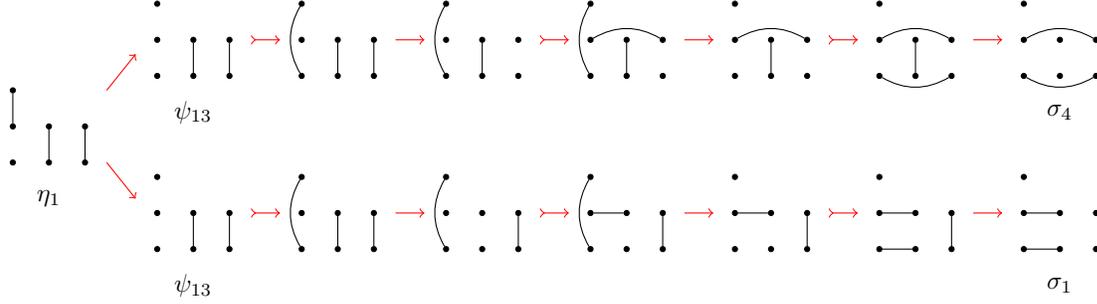
Fig. 11: Only two possible $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_1, and end at a critical 1-simplex.

By Equation (2.2), the multiplicity of γ_1,

$$m(\gamma_1) = (-\langle \beta_0, \alpha_0 \rangle \langle \beta_0, \alpha_1 \rangle)(-\langle \beta_1, \alpha_1 \rangle \langle \beta_1, \alpha_2 \rangle)(-\langle \beta_2, \alpha_2 \rangle \langle \beta_2, \alpha_3 \rangle)$$
$$= (-(+1)(+1))(-(+1)(+1))(-(+1)(-1)) = +1.$$

Let γ_2 be the unique $\mathcal{M}^*$-path that starts from ψ_{13}, and ends at σ_1 (see Fig. 10 and Fig. 11), i.e.,

$$\gamma_2 : \psi_{13} = \alpha_0 \xrightarrow{(+1)} \beta_0 \xrightarrow{(-1)} \alpha_1 \xrightarrow{(+1)} \beta_1 \xrightarrow{(+1)} \alpha_2 \xrightarrow{(+1)} \beta_2 \xrightarrow{(-1)} \alpha_3 = \sigma_1, \quad \text{where}$$

$$\alpha_0 = [v_2^{(1)} v_2^{(2)}, v_3^{(1)} v_3^{(2)}] = \psi_{13}, \qquad \beta_0 = [v_1^{(1)} v_1^{(3)}, v_2^{(1)} v_2^{(2)}, v_3^{(1)} v_3^{(2)}]$$
$$\alpha_1 = [v_1^{(1)} v_1^{(3)}, v_3^{(1)} v_3^{(2)}], \qquad \beta_1 = [v_1^{(1)} v_1^{(3)}, v_3^{(1)} v_3^{(2)}, v_1^{(2)} v_2^{(2)}]$$
$$\alpha_2 = [v_3^{(1)} v_3^{(2)}, v_1^{(2)} v_2^{(2)}], \qquad \beta_2 = [v_1^{(1)} v_2^{(1)}, v_3^{(1)} v_3^{(2)}, v_1^{(2)} v_2^{(2)}]$$
$$\alpha_3 = [v_1^{(1)} v_2^{(1)}, v_1^{(2)} v_2^{(2)}] = \sigma_1.$$

By Equation (2.2), the multiplicity of γ_2,

$$m(\gamma_2) = (-\langle \beta_0, \alpha_0 \rangle \langle \beta_0, \alpha_1 \rangle)(-\langle \beta_1, \alpha_1 \rangle \langle \beta_1, \alpha_2 \rangle)(-\langle \beta_2, \alpha_2 \rangle \langle \beta_2, \alpha_3 \rangle)$$
$$= (-(+1)(-1))(-(+1)(+1))(-(+1)(-1)) = -1.$$

Therefore, from Equation (2.3) and Equation (2.4),

$$\tilde{\partial}_2(\eta_1) = (\langle \eta_1, \psi_{13} \rangle \cdot m(\gamma_1)) \cdot \sigma_4 + (\langle \eta_1, \psi_{13} \rangle \cdot m(\gamma_2)) \cdot \sigma_1$$
$$= ((+1)(+1))) \cdot \sigma_4 + ((+1)(-1)) \cdot \sigma_1 = -\sigma_1 + \sigma_4.$$

Considering all possible $\mathcal{M}^*$-paths that start from a 1-simplex contained in a critical 2-simplex, and end at a critical 1-simplex (refer to Appendix A), we may compute the images of all twenty-four critical 2-simplices under the boundary operator $\tilde{\partial}_2$. These are all listed in Tab. 1. We describe an algorithmic scheme to compute the images of critical 2-simplices under $\tilde{\partial}_2$ in Appendix B, and the computation of boundaries of two more critical 2-simplices following this scheme is also discussed there.

η	$\tilde{\partial}_2(\eta)$	η	$\tilde{\partial}_2(\eta)$	η	$\tilde{\partial}_2(\eta)$	η	$\tilde{\partial}_2(\eta)$
	$(-\sigma_1+\sigma_1+\sigma_2-\sigma_3+\sigma_4-\sigma_4)=\sigma_2-\sigma_3$		$\sigma_1-\sigma_2-\sigma_3$		$\sigma_1-\sigma_3+\sigma_4$		$\sigma_2+\sigma_3-\sigma_4$
	$(\sigma_1-\sigma_2+\sigma_2+\sigma_3-\sigma_3-\sigma_4)=\sigma_1-\sigma_4$		$-\sigma_1+\sigma_2-\sigma_4$		$-\sigma_1+\sigma_4$		$-\sigma_2+\sigma_3$
	$(-\sigma_1+\sigma_2+\sigma_3+\sigma_4-\sigma_4)=-\sigma_1+\sigma_2+\sigma_3$		$-\sigma_1+\sigma_3-\sigma_4$		$-\sigma_1+\sigma_2+\sigma_3$		$\sigma_1-\sigma_2+\sigma_4$
	$(-\sigma_1+\sigma_1-\sigma_2-\sigma_3+\sigma_4)=-\sigma_2-\sigma_3+\sigma_4$		$-\sigma_2-\sigma_3+\sigma_4$		$\sigma_2-\sigma_3$		$\sigma_1-\sigma_4$
	$(\sigma_1-\sigma_2+\sigma_3-\sigma_3+\sigma_4)=\sigma_1-\sigma_2+\sigma_4$		$-\sigma_1+\sigma_4$		$-\sigma_1+\sigma_2-\sigma_4$		$\sigma_1-\sigma_2-\sigma_3$
	$(-\sigma_1-\sigma_2+\sigma_2+\sigma_3-\sigma_4)=-\sigma_1+\sigma_3-\sigma_4$		$-\sigma_2+\sigma_3$		$\sigma_2+\sigma_3-\sigma_4$		$\sigma_1-\sigma_3+\sigma_4$

Tab. 1: Images of all critical 2-simplices under the boundary operator $\tilde{\partial}_2$.

4.2.3 First and second homology groups of M_7

We have $\ker(\tilde{\partial}_1) = \tilde{C}_1(M_7) = \left\langle \sigma_1, \sigma_2, \sigma_3, \sigma_4 \right\rangle$, and from Tab. 1, we conclude that $\operatorname{im}(\tilde{\partial}_2)$ is generated by $\sigma_1 - \sigma_4, \sigma_2 - \sigma_3, \sigma_1 - \sigma_2 - \sigma_3, \sigma_1 - \sigma_2 + \sigma_4, \sigma_1 - \sigma_3 + \sigma_4$, and $\sigma_2 + \sigma_3 - \sigma_4$. Thus,

$$
\begin{aligned}
& H_1(M_7) \\
\cong\ & \ker(\tilde{\partial}_1)\big/\operatorname{im}(\tilde{\partial}_2) \\
\cong\ & \left\langle \sigma_1, \sigma_2, \sigma_3, \sigma_4 \right\rangle\big/\left\langle \sigma_1 - \sigma_4, \sigma_2 - \sigma_3, \sigma_1 - \sigma_2 - \sigma_3, \sigma_1 - \sigma_2 + \sigma_4, \sigma_1 - \sigma_3 + \sigma_4, \sigma_2 + \sigma_3 - \sigma_4 \right\rangle \\
\cong\ & \left\langle \sigma_1, \sigma_2 \right\rangle\big/\left\langle \sigma_1 - 2\sigma_2, 2\sigma_1 - \sigma_2 \right\rangle \\
\cong\ & \left\langle \sigma_1 \right\rangle\big/\left\langle 3\sigma_1 \right\rangle \\
\cong\ & \mathbb{Z}/3\mathbb{Z} = \mathbb{Z}_3.
\end{aligned}
$$

Finally, from Tab. 1, we may verify that out of images of all 24 generators of $\tilde{C}_2(M_7)$ under the boundary operator $\tilde{\partial}_2$, there are exactly four linearly independent ones. Therefore, $\operatorname{rank}(\ker(\tilde{\partial}_2)) = 24 - 4 = 20$. Hence, $H_2(M_7) \cong \mathbb{Z}^{20}$. This concludes the proof of Theorem 1.3.

Remark 4.4. After computing $H_1(M_7)$, one may also compute the second homology group of M_7 using the Euler characteristic of M_7 as follows. If f_i is the number of i-dimensional simplices of M_7, then from Equation (2.1), $f_0 = 21$, $f_1 = 105$, and $f_2 = 105$ (with $f_i = 0$, for all $i \geq 3$). Thus, $\chi(M_7) = \sum_{i \geq 0}(-1)^i f_i = 21 - 105 + 105 = 21$. If b_i is the i-the Betti number of M_7, then $b_0 = 1$, $b_1 = 0$, and $b_i = 0$ for $i \geq 3$ (since M_7 is a connected 2-dimensional complex with $H_1(M_7) = \mathbb{Z}_3$). Thus, $b_2 = \chi(M_7) - b_0 = 20$. Since $H_2(M_7)$ is torsion-free, we get $H_2(M_7) = \mathbb{Z}^{20}$.

4.3 Further augmentation of the near-optimal gradient vector field

We note that $\mathcal{M}^*$ is a highly efficient (near-optimal) gradient vector field on M_7 as, from Observation 4.2, it follows that

  (i) out of 105 simplices of dimension 2, only 24 are critical (with the second Betti number being 20),

  (ii) out of 105 simplices of dimension 1, only four are critical (with the first Betti number being 0),

 (iii) out of 21 simplices of dimension 0, only one is critical (with the zeroth Betti number being 1).

However, we may augment $\mathcal{M}^*$ even further to one that satisfies Theorem 1.4.

Proof of Theorem 1.4: From Fig. 11 and Fig. 25 (Appendix B), we observe that there is

  (i) a unique $\mathcal{M}^*$-path that starts from a 1-simplex (viz., ψ_{13}) contained in the critical 2-simplex η_1, and ends at the critical 1-simplex σ_4,

  (ii) a unique $\mathcal{M}^*$-path that starts from a 1-simplex (viz., ψ_{22}) contained in the critical 2-simplex η_2, and ends at the critical 1-simplex σ_3,

as shown in Fig. 12. Moreover, we also observe that there is

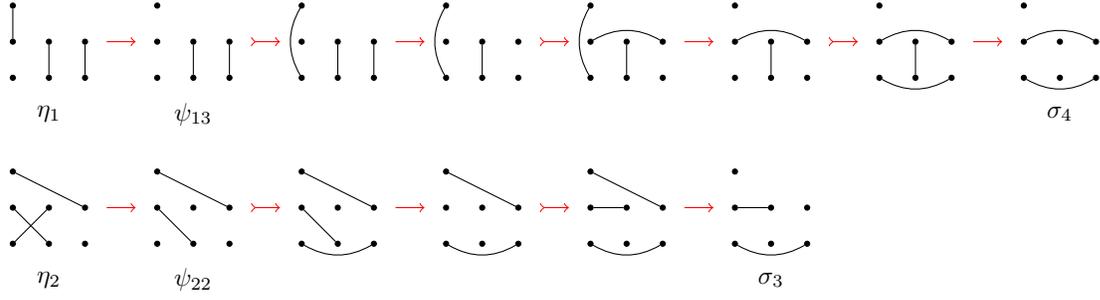 (i) no $\mathcal{M}^*$-path that starts from a 1-simplex contained in η_1, and ends at σ_3,

Fig. 12: Unique $\mathcal{M}^*$-paths starting from a 1-simplex contained in η_1 and η_2, and ending at σ_4 and σ_3, respectively.

(ii) no $\mathcal{M}^*$-path that starts from a 1-simplex contained in η_2, and ends at σ_4.

Thus, the requirements of Theorem 2.8 are satisfied. This allows us to apply the technique of cancellation of a critical pair by reversing paths (Theorem 2.7) twice, and end up with a gradient vector field $\mathcal{M}^{**}$ from $\mathcal{M}^*$ such that η_1, σ_4, and η_2, σ_3 are not critical with respect to $\mathcal{M}^{**}$, while the criticality of all other simplices remains unchanged. Thus, with respect to $\mathcal{M}^{**}$, there are $24 - 2 = 22$ critical 2-simplices, two critical 1-simplices (viz., σ_1 and σ_2), and one critical 0-simplex (viz., ξ). □

By Theorem 2.4, it follows that M_7 is homotopy equivalent to a CW complex with one 0-cell, two 1-cells, and twenty-two 2-cells.

5   Conclusion

As $H_1(M_7)$ is nontrivial, with respect to any gradient vector field on M_7, there is at least one critical 1-simplex. As a consequence, from the Morse inequalities (Theorem 2.6), it follows that with respect to any gradient vector field on M_7, there are at least 21 critical 2-simplices.

We note that (from Appendix A) there is an $\mathcal{M}^*$-path from a face of a critical (with respect to $\mathcal{M}^*$) 2-simplex, say η, to the critical 1-simplex σ_1 if and only if there is an $\mathcal{M}^*$-path from a face of η to σ_4. Suppose we first cancel the critical pair $\eta^{(2)}, \sigma_i^{(1)}$, where $i \in \{1, 4\}$, by reversing the unique gradient path γ from a face of η to σ_i (by Theorem 2.7). Then in addition to an $\mathcal{M}^*$-path (if one exists) from a face of another critical 2-simplex η' to σ_{5-i}, there is also a *new* gradient path (with respect to the augmented gradient vector field) from a face of η' to σ_{5-i} that passes through reversed γ. Thus, the uniqueness of possible gradient paths (with respect to the augmented gradient vector field) from a face of any remaining critical 2-simplex to σ_{5-i}, is no longer viable. So, out of σ_1 and σ_4, *at most one* can be cancelled. Likewise, at most one of σ_2 and σ_3 can be cancelled. Therefore, starting with $\mathcal{M}^*$, the technique of cancellation of critical pairs can be applied *at most* twice. This motivates us to conjecture that the augmented gradient vector field $\mathcal{M}^{**}$, which is obtained after applying the cancellation theorem twice, is indeed an optimal one.

Conjecture 5.1. *A gradient vector field on M_7 with respect to which there are 22 critical 2-simplices, two critical 1-simplices, and one critical 0-simplex, is an optimal one.*

As mentioned before, in order to compute the (discrete) Morse homology groups of M_7 in an efficient manner, we constructed and worked with the near-optimal gradient vector field $\mathcal{M}^*$ that not only admits

a low number of critical simplices of each dimension, but also, more importantly, allows an efficient (and human-comprehensible) enumeration of all possible gradient paths of interest. It would be another problem of interest to determine how "far from random" the constructed gradient vector fields are.

Problem 5.2. *Compare the efficacy of the gradient vector field $\mathcal{M}^*$ (and the augmented one, $\mathcal{M}^{**}$) with those found by random methods as described by Adiprasito et al. (2017), Benedetti and Lutz (2014), Benedetti et al. (2023).*

We also expect the techniques developed in this article to be useful for taking up problems of exploring the topology of higher dimensional matching complexes. As discussed in this article, one potential approach consists of the following two steps in order.

1. Extend the initial gradient vector field constructed in Section 3, which is optimal up to dimension $\nu_n - 1$, to a sufficiently good (optimal or near-optimal) one. Some augmentation tricks used in this article may be adapted for higher dimensions. For example, the augmentation used in Example 3.3 is applicable to any M_{2n}, and the augmentation used in Subsection 4.1 may be applied to any M_{3n+1}.

2. Using the efficient extended gradient vector field, (perhaps with the aid of computer implementations) compute the homology groups.

Appendix A $\mathcal{M}^*$-paths ending at a critical 1-simplex

Let us consider the critical 1-simplex $\sigma_1 = \{v_1^{(1)}v_2^{(1)}, v_1^{(2)}v_2^{(2)}\}$ as shown below.

If $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)}, \alpha_{k+1}^{(1)}$ is an $\mathcal{M}^*$-path ending at σ_1 (i.e., $\alpha_{k+1} = \sigma_1$), then β_k is one of the following three 2-simplices.
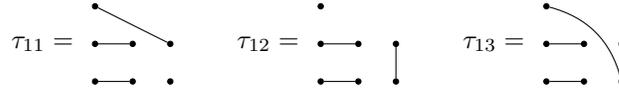
In Fig. 13, we describe *all possible* (following the diagram from right to left would make it clear) $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{11}, \alpha_{k+1}^{(1)} = \sigma_1,$$

where α_0 is contained in a critical 2-simplex η (possible choices of $\eta^{(2)}$ are also included in the leftmost column of the figure).
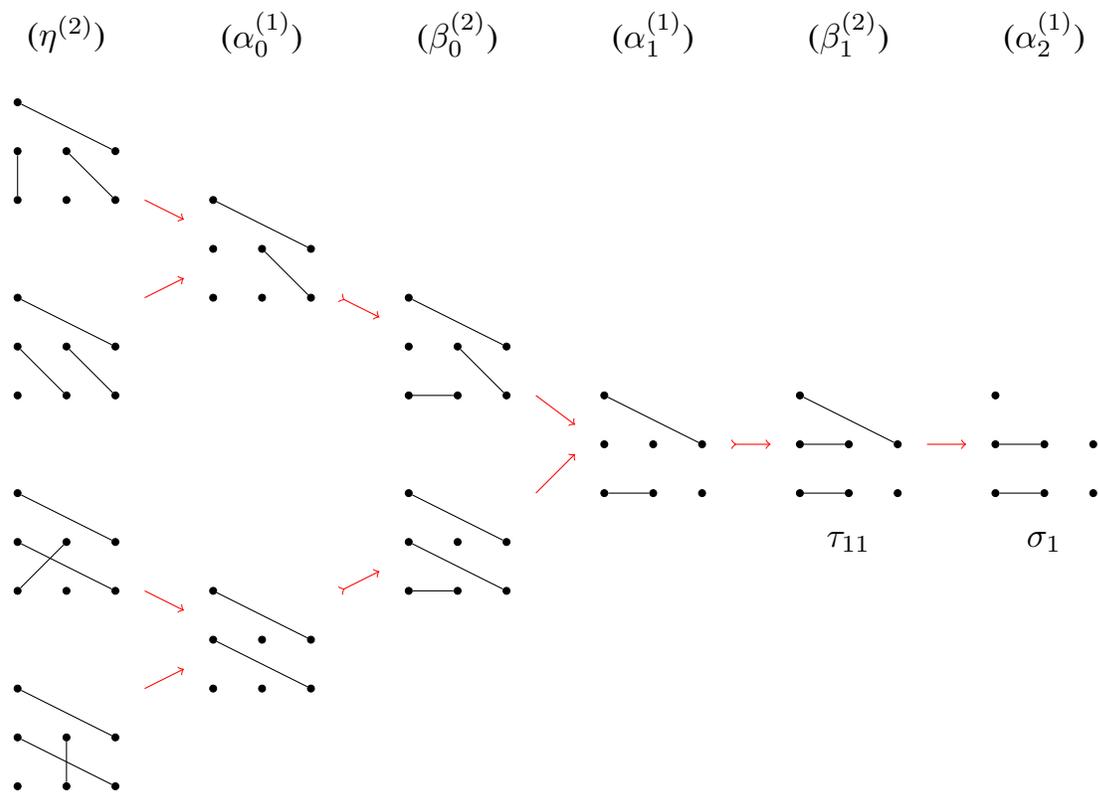
Similarly, in Fig. 14, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{12}, \alpha_{k+1}^{(1)} = \sigma_1,$$

and in Fig. 15, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{13}, \alpha_{k+1}^{(1)} = \sigma_1,$$

where α_0 is contained in a critical 2-simplex η.

$(\eta^{(2)})$ $(\alpha_0^{(1)})$ $(\beta_0^{(2)})$ $(\alpha_1^{(1)})$ $(\beta_1^{(2)})$ $(\alpha_2^{(1)})$

τ_{11} σ_1

Fig. 13: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{11}, \alpha_{k+1}^{(1)} = \sigma_1$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
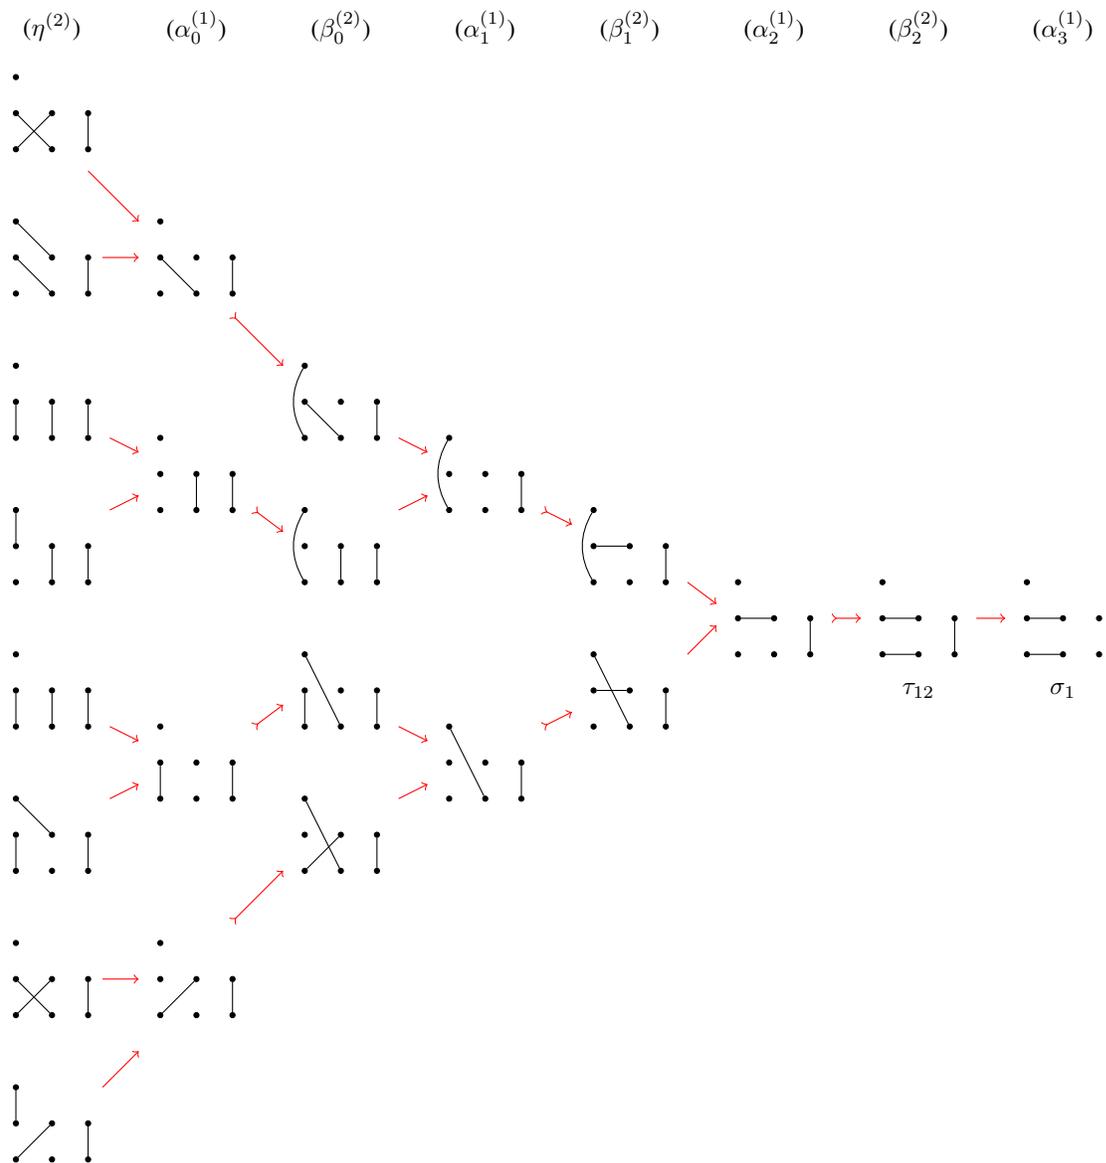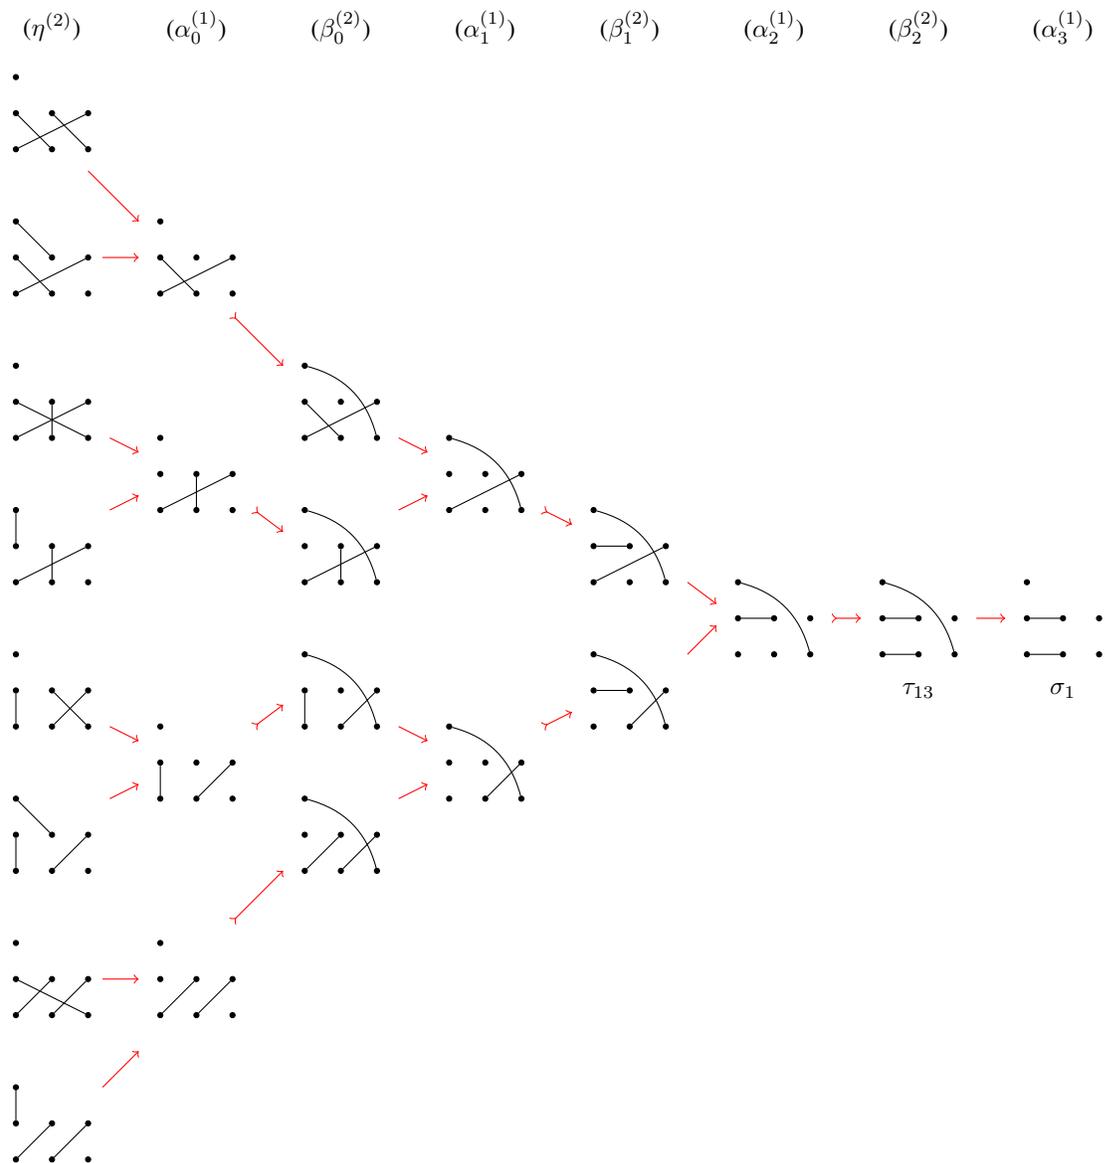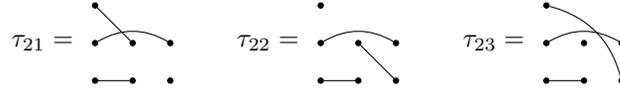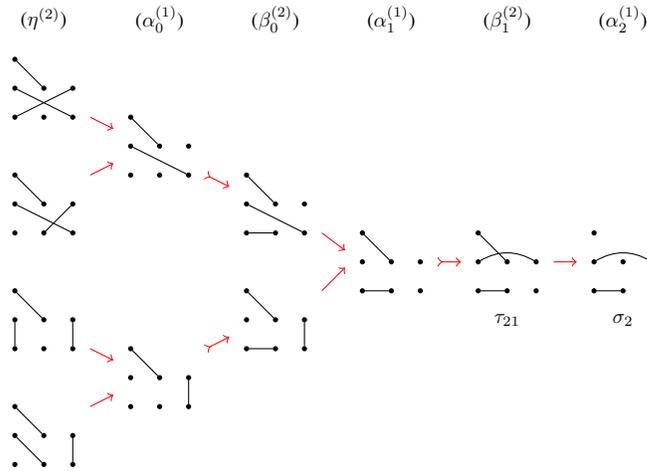
Fig. 14: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{12}, \alpha_{k+1}^{(1)} = \sigma_1$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
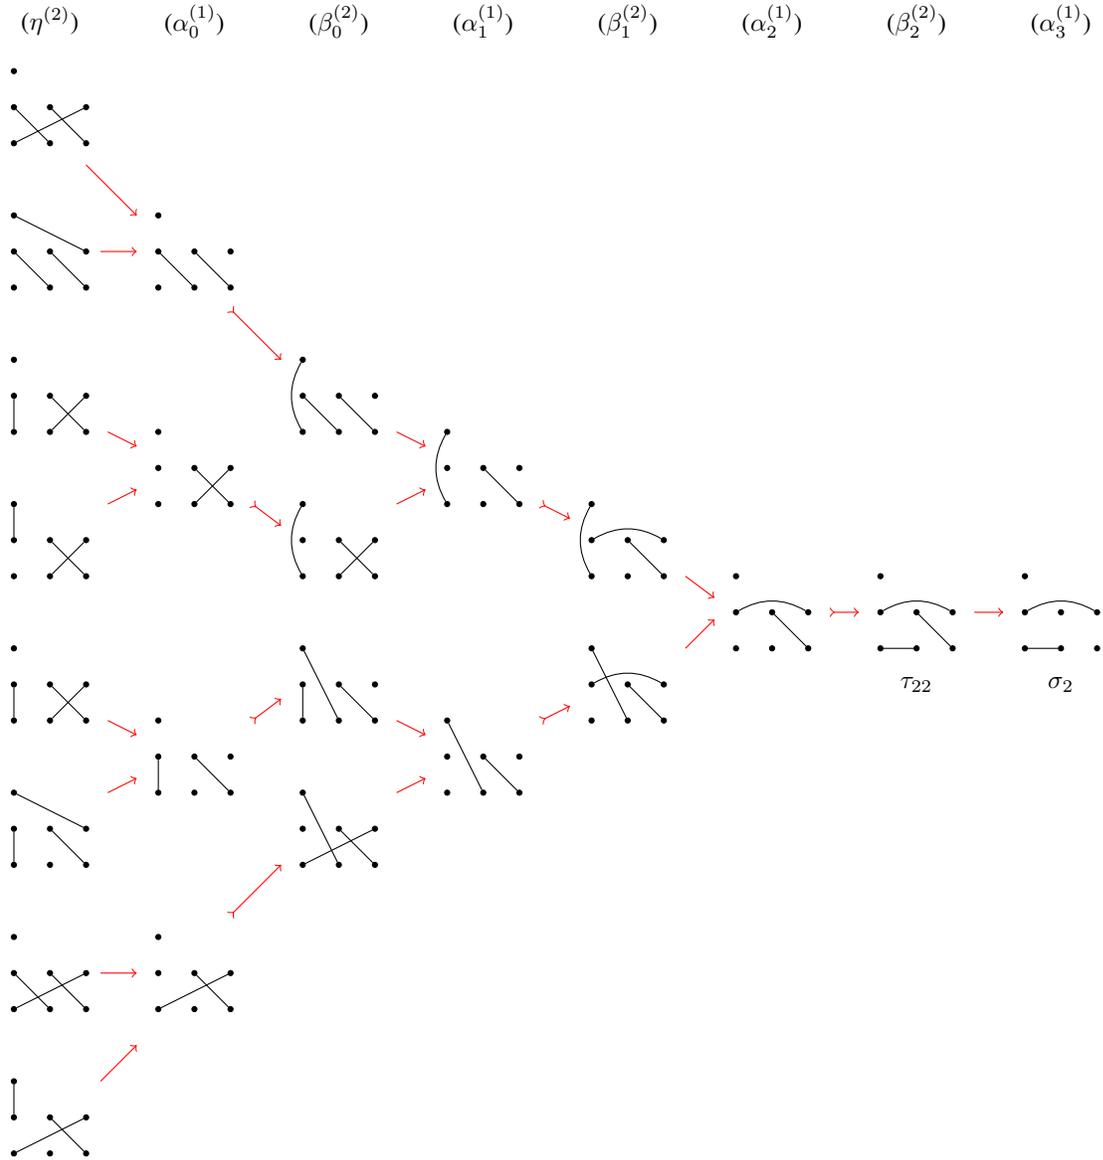
Fig. 15: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{13}, \alpha_{k+1}^{(1)} = \sigma_1$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Now, let us consider the critical 1-simplex $\sigma_2 = \{v_1^{(1)} v_2^{(1)}, v_1^{(2)} v_3^{(2)}\}$ as shown below.

If $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)}, \alpha_{k+1}^{(1)}$ is an $\mathcal{M}^*$-path ending at σ_2 (i.e., $\alpha_{k+1} = \sigma_2$), then β_k is one of the following three 2-simplices.

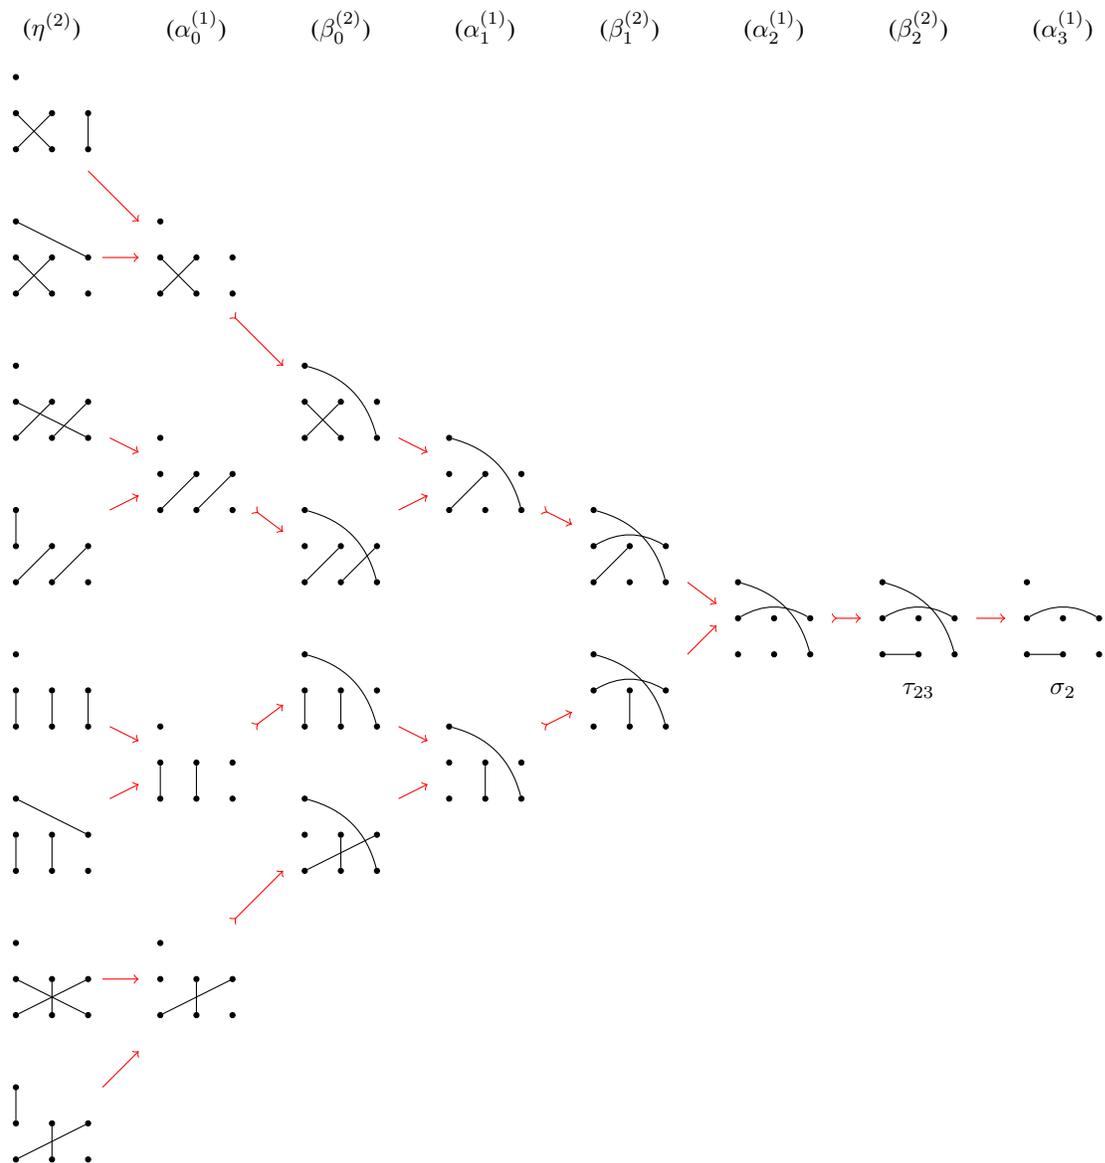In Fig. 16, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{21}, \alpha_{k+1}^{(1)} = \sigma_2$$

in Fig. 17, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{22}, \alpha_{k+1}^{(1)} = \sigma_2$$

and in Fig. 18, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{23}, \alpha_{k+1}^{(1)} = \sigma_2,$$

where α_0 is contained in a critical 2-simplex η.

Fig. 16: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{21}, \alpha_{k+1}^{(1)} = \sigma_2$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
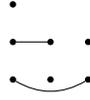
$(\eta^{(2)})$ $(\alpha_0^{(1)})$ $(\beta_0^{(2)})$ $(\alpha_1^{(1)})$ $(\beta_1^{(2)})$ $(\alpha_2^{(1)})$ $(\beta_2^{(2)})$ $(\alpha_3^{(1)})$

τ_{22} σ_2

Fig. 17: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}$, $\beta_0^{(2)}$, $\alpha_1^{(1)}$, $\beta_1^{(2)}$, ..., $\alpha_k^{(1)}$, $\beta_k^{(2)} = \tau_{22}$, $\alpha_{k+1}^{(1)} = \sigma_2$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
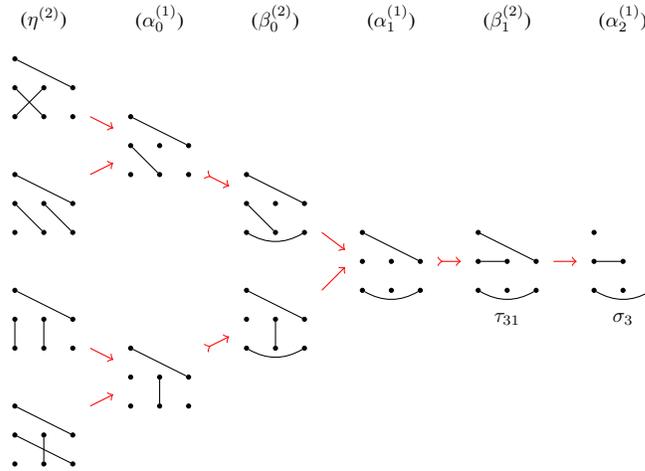
Fig. 18: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{23}, \alpha_{k+1}^{(1)} = \sigma_2$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Next, let us consider the critical 1-simplex $\sigma_3 = \{v_1^{(1)}v_3^{(1)}, v_1^{(2)}v_2^{(2)}\}$ as shown below.

If $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)}, \alpha_{k+1}^{(1)}$ is an $\mathcal{M}^*$-path ending at σ_3 (i.e., $\alpha_{k+1} = \sigma_3$), then β_k is one of the following three 2-simplices.
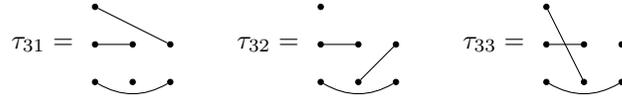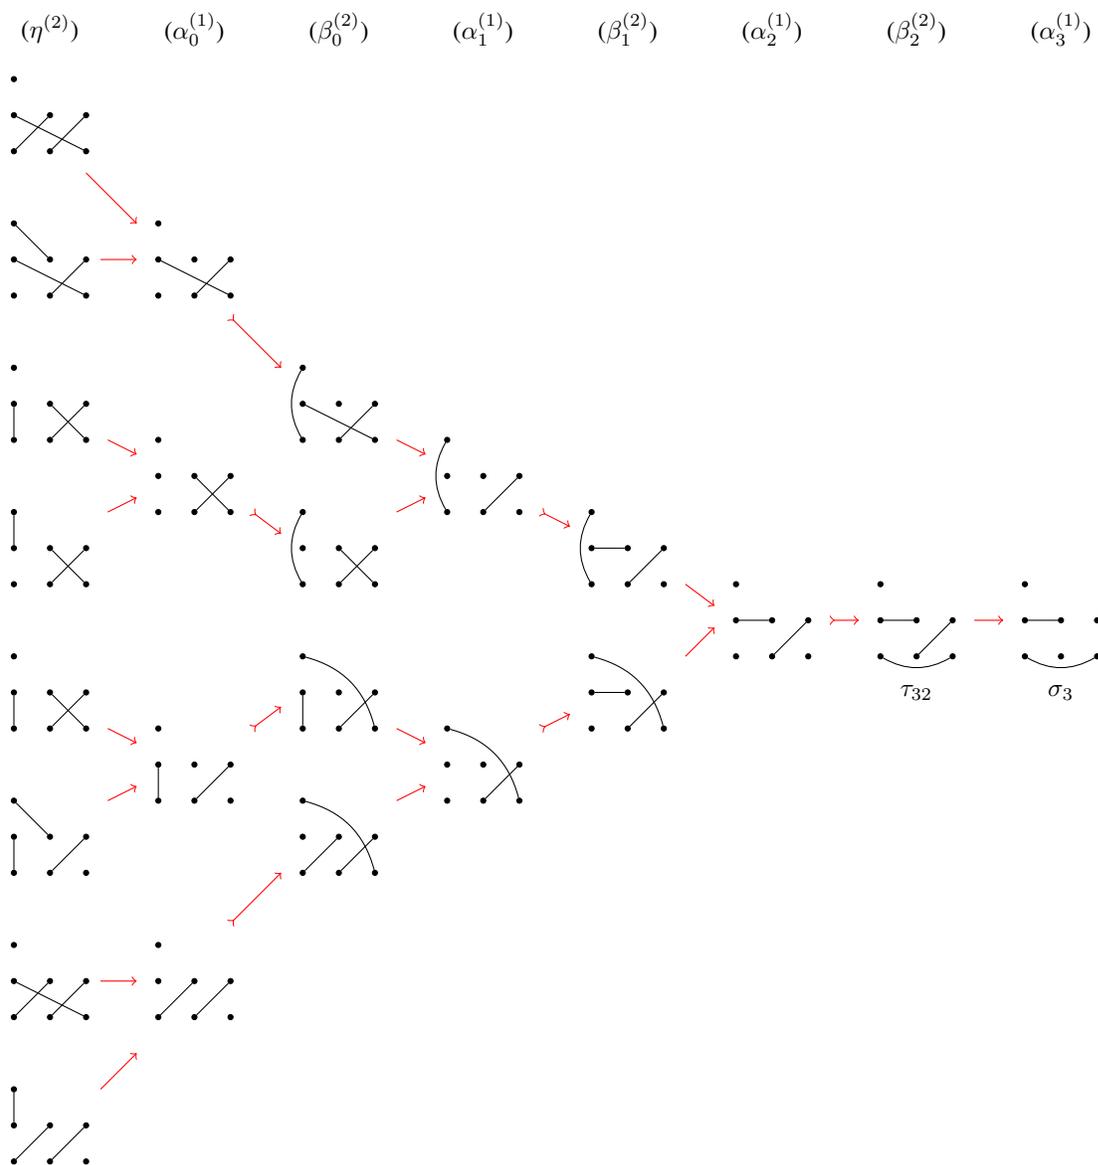

In Fig. 19, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{31}, \alpha_{k+1}^{(1)} = \sigma_3,$$

in Fig. 20, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{32}, \alpha_{k+1}^{(1)} = \sigma_3,$$

and in Fig. 21, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{33}, \alpha_{k+1}^{(1)} = \sigma_3,$$

where α_0 is contained in a critical 2-simplex η.

Fig. 19: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{31}, \alpha_{k+1}^{(1)} = \sigma_3$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
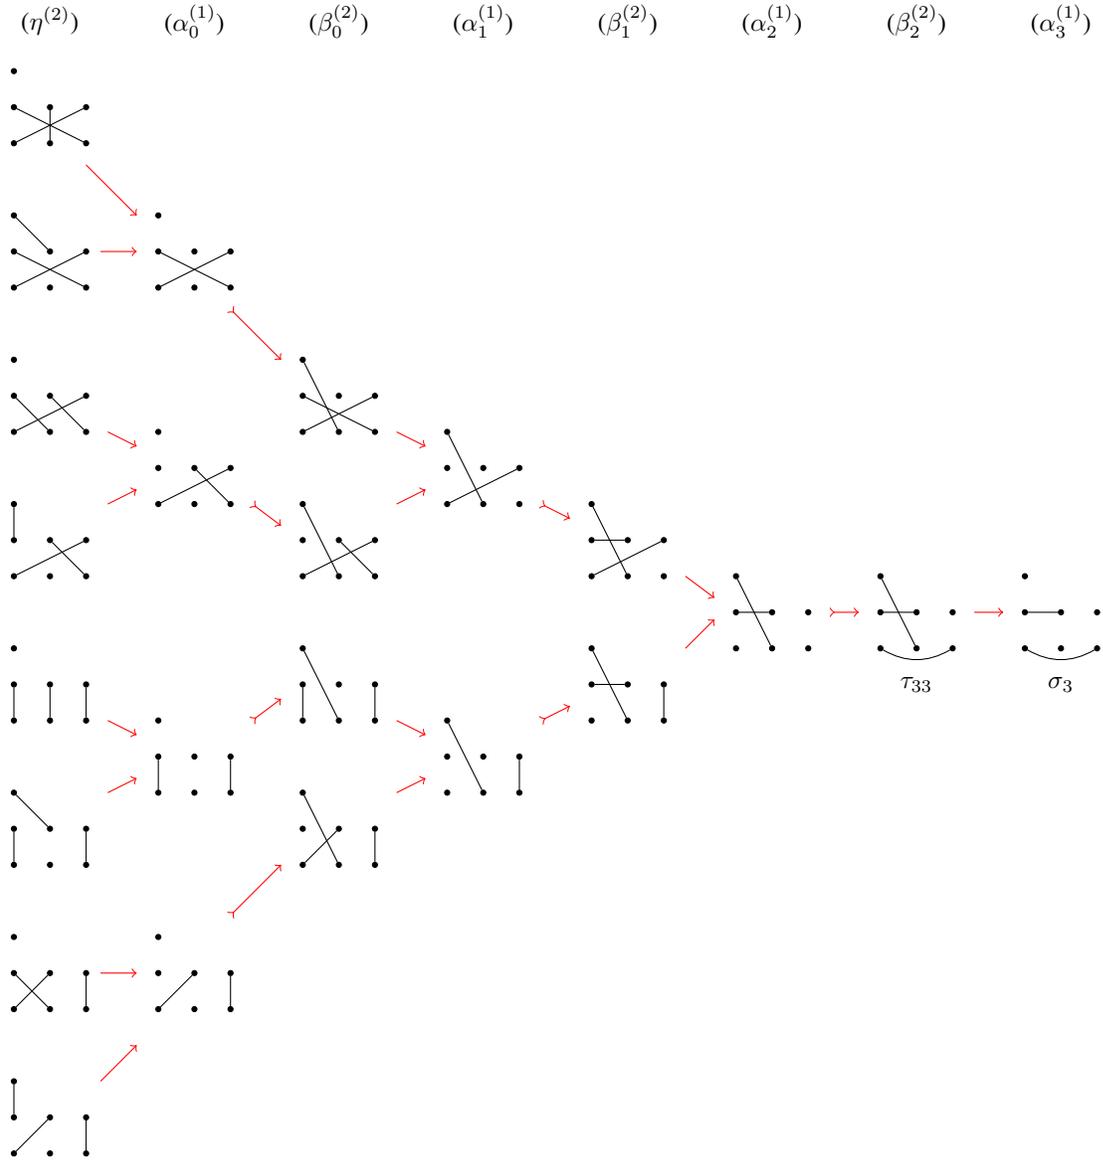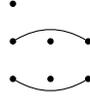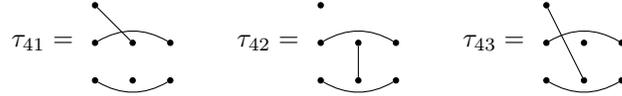
Fig. 20: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{32}, \alpha_{k+1}^{(1)} = \sigma_3$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Fig. 21: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}$, $\beta_0^{(2)}$, $\alpha_1^{(1)}$, $\beta_1^{(2)}$, ..., $\alpha_k^{(1)}$, $\beta_k^{(2)} = \tau_{33}$, $\alpha_{k+1}^{(1)} = \sigma_3$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Finally, let us consider the critical 1-simplex $\sigma_4 = \{v_1^{(1)}v_3^{(1)}, v_1^{(2)}v_3^{(2)}\}$ as shown below.

If $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)}, \alpha_{k+1}^{(1)}$ is an $\mathcal{M}^*$-path ending at σ_4 (i.e., $\alpha_{k+1} = \sigma_4$), then β_k is one of the following three 2-simplices.

$$\tau_{41} = \qquad \tau_{42} = \qquad \tau_{43} =$$

In Fig. 22, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{41}, \alpha_{k+1}^{(1)} = \sigma_4,$$
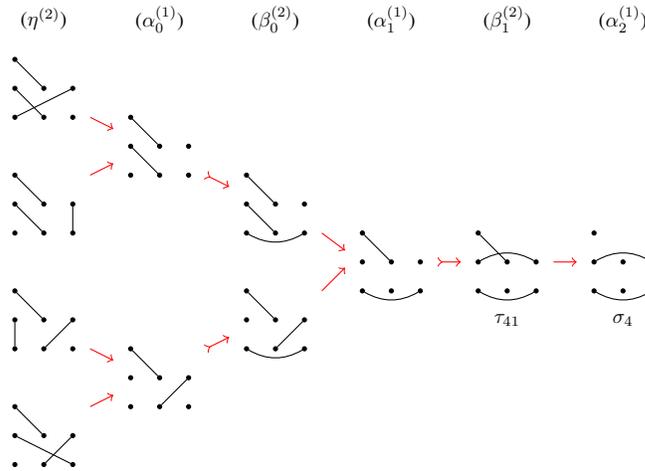
in Fig. 23, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{42}, \alpha_{k+1}^{(1)} = \sigma_4,$$

and in Fig. 24, we describe all possible $\mathcal{M}^*$-paths of the form

$$\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{43}, \alpha_{k+1}^{(1)} = \sigma_4,$$

where α_0 is contained in a critical 2-simplex η.

$$(\eta^{(2)}) \qquad (\alpha_0^{(1)}) \qquad (\beta_0^{(2)}) \qquad (\alpha_1^{(1)}) \qquad (\beta_1^{(2)}) \qquad (\alpha_2^{(1)})$$
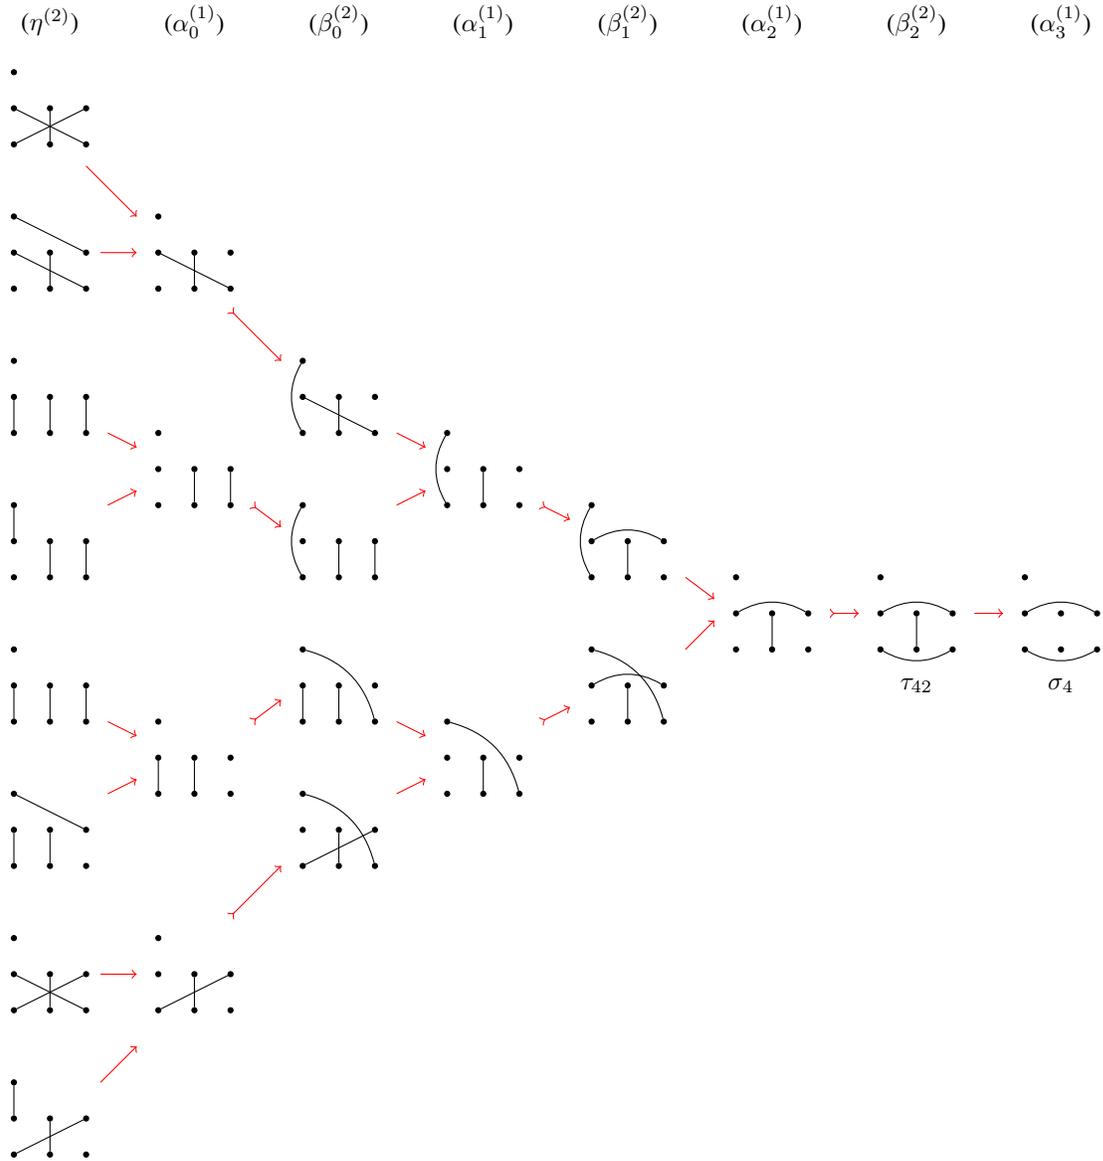
Fig. 22: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{41}, \alpha_{k+1}^{(1)} = \sigma_4$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.
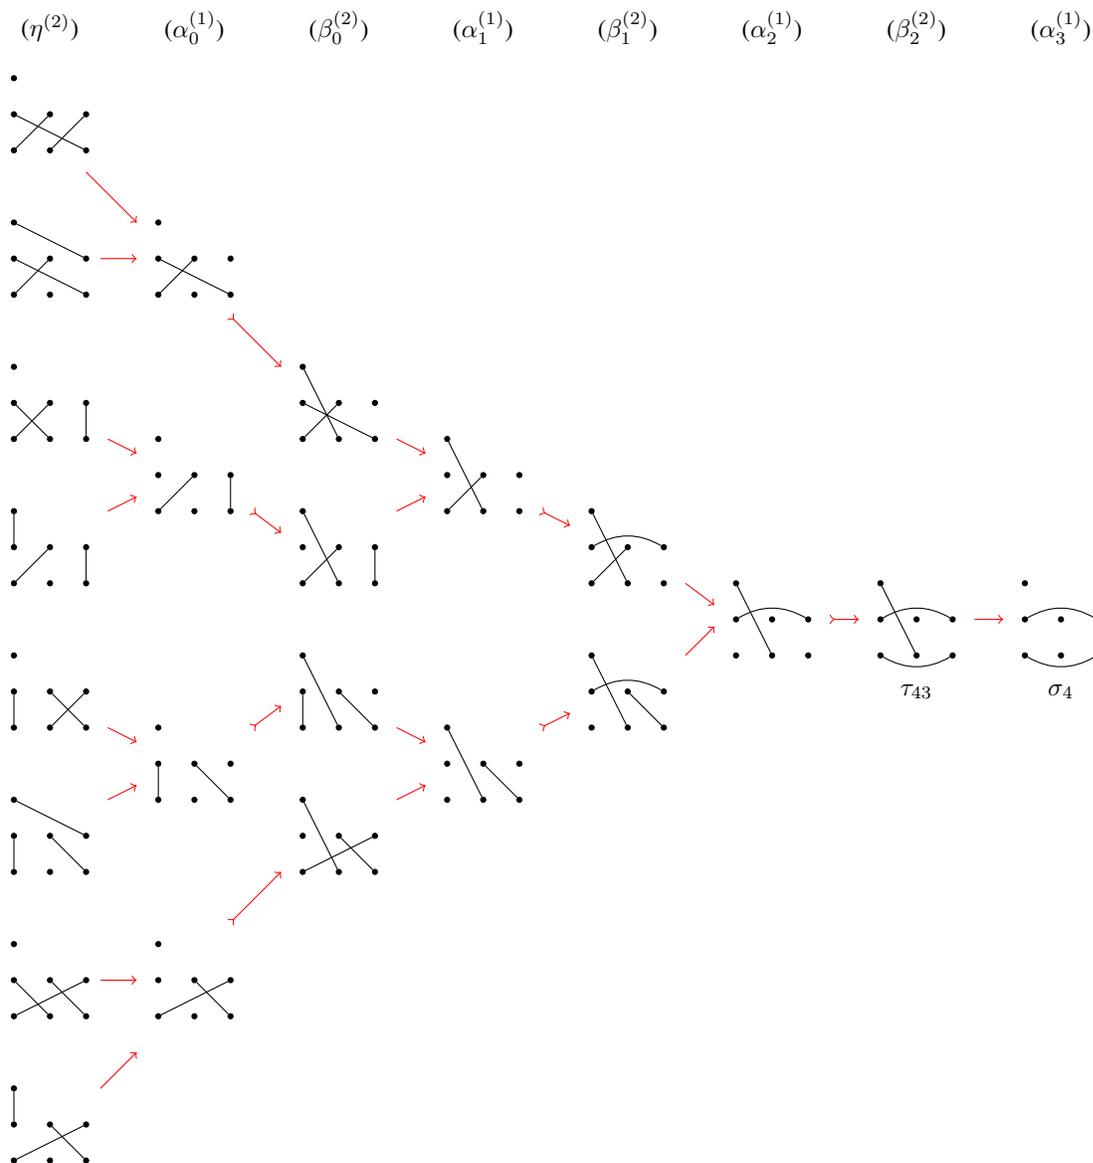
Fig. 23: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{42}, \alpha_{k+1}^{(1)} = \sigma_4$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Fig. 24: All possible $\mathcal{M}^*$-paths of the form $\alpha_0^{(1)}, \beta_0^{(2)}, \alpha_1^{(1)}, \beta_1^{(2)}, \ldots, \alpha_k^{(1)}, \beta_k^{(2)} = \tau_{43}, \alpha_{k+1}^{(1)} = \sigma_4$, with $\alpha_0 \subsetneq \eta$, where η is a critical 2-simplex.

Appendix B   A scheme to compute the boundaries of critical 2-simplices

First, in each figure in Appendix A, above each arrow connecting $\beta_i^{(2)}$ and $\alpha_j^{(1)}$ (with $0 \le i \le j \le i+1$), we write the incidence number $\langle \beta_i, \alpha_j \rangle$. Also, above each arrow from the critical 2-simplex η (leftmost

column) to the 1-simplex α_0, we write the incidence number $\langle \eta, \alpha_0 \rangle$, as shown below.

$$\eta \xrightarrow{\langle \eta, \alpha_0 \rangle} \alpha_0 \xrightarrowtail{\langle \beta_0, \alpha_0 \rangle} \beta_0 \xrightarrow{\langle \beta_0, \alpha_1 \rangle} \alpha_1 \xrightarrowtail{\langle \beta_1, \alpha_1 \rangle} \beta_1 \xrightarrow{\langle \beta_1, \alpha_2 \rangle} \cdots \xrightarrow{\langle \beta_{k-1}, \alpha_k \rangle} \alpha_k \xrightarrowtail{\langle \beta_k, \alpha_k \rangle} \beta_k \xrightarrow{\langle \beta_k, \alpha_{k+1} \rangle} \alpha_{k+1}$$

Next, let us choose a particular critical 2-simplex η_0. We describe a sequence of algorithmic steps to compute $\tilde{\partial}_2(\eta_0)$ below.

Step 1: Pick a path from η_0 to a critical 1-simplex, say σ_i, in a figure of Appendix A.

Step 2: Count the number of -1's appearing above the arrows in the path. Let the count be r.

Step 3: Count the number of non-critical 2-simplices in the path (i.e., the number of $\beta_i^{(2)}$s). Let the count be s. The contribution of this specific path to $\tilde{\partial}_2(\eta_0)$ is $(-1)^{r+s}\sigma_i$.

Step 4: Compute the contribution of each possible path from η_0 to a critical 1-simplex in all the figures of Appendix A.

Step 5: Sum of the contributions of all possible paths, from η_0 to a critical 1-simplex, is $\tilde{\partial}_2(\eta_0)$.

Let us consider a couple of critical 2-simplices and compute their images under the boundary operator $\tilde{\partial}_2$ following the scheme above in the examples below.

Example B.1. Let us consider the critical 2-simplex $\eta_2 := \{v_1^{(1)}v_2^{(2)}, v_2^{(1)}v_1^{(2)}, v_3^{(2)}v_1^{(3)}\}$ as shown below.

We observe that there are exactly two $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_2, and end at a critical 1-simplex (Fig. 18 and Fig. 19 in Appendix A) as shown in Fig. 25 (along with the incidence numbers).
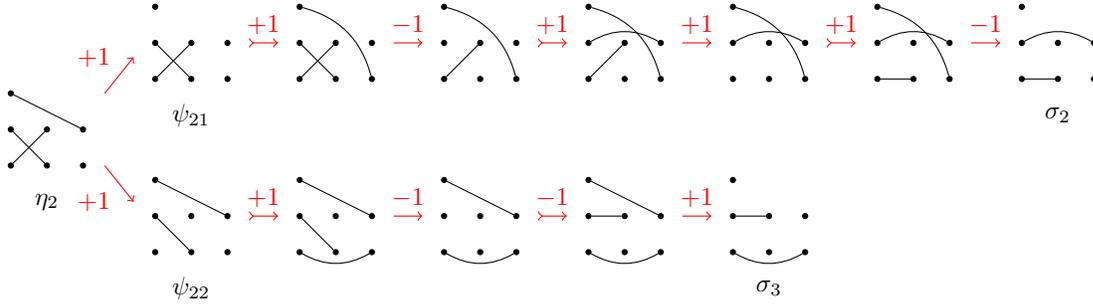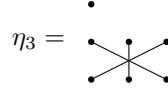

Fig. 25: Only two possible $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_2, and end at a critical 1-simplex.

- Let P_1 be the path from η_2 to σ_2 via ψ_{21} in Fig. 25. There are two -1's appearing in P_1 and there are three non-critical 2-simplices in P_1. Thus, P_1 contributes $(-1)^{2+3}\sigma_2 = -\sigma_2$ to $\tilde{\partial}_2(\eta_2)$.

- Let P_2 be the path from η_2 to σ_3 via ψ_{22} in Fig. 25. There are two -1's appearing in P_2 and there are two non-critical 2-simplices in P_2. Thus, P_2 contributes $(-1)^{2+2}\sigma_3 = \sigma_3$ to $\tilde{\partial}_2(\eta_2)$.

Therefore, $\tilde{\partial}_2(\eta_2) = -\sigma_2 + \sigma_3$.

Example B.2. Let us consider the critical 2-simplex $\eta_3 := \{v_1^{(1)}v_3^{(2)}, v_2^{(1)}v_2^{(2)}, v_3^{(1)}v_1^{(2)}\}$ as shown below.

$$\eta_3 = \qquad$$

We observe that there are five $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_3, and end at a critical 1-simplex (Fig. 15, Fig. 18, Fig. 21, and Fig. 23 in Appendix A) as shown in Fig. 26 (along with the incidence numbers).
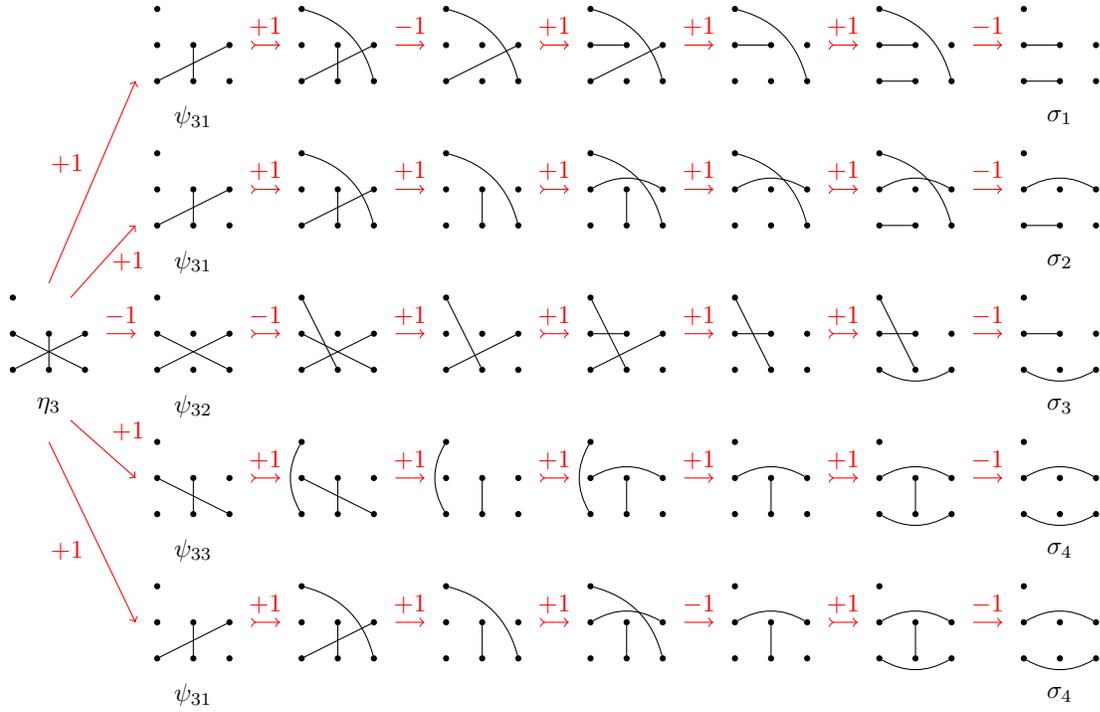

Fig. 26: All five $\mathcal{M}^*$-paths that start from a 1-simplex contained in η_3, and end at a critical 1-simplex.

- Let P_1 be the path from η_3 to σ_1 via ψ_{31} in Fig. 26. There are two -1's appearing in P_1 and there are three non-critical 2-simplices in P_1. Thus, P_1 contributes $(-1)^{2+3}\sigma_1 = -\sigma_1$ to $\tilde{\partial}_2(\eta_3)$.

- Let P_2 be the path from η_3 to σ_2 via ψ_{31} in Fig. 26. There is only one -1 appearing in P_2 and there are three non-critical 2-simplices in P_2. Thus, P_2 contributes $(-1)^{1+3}\sigma_2 = \sigma_2$ to $\tilde{\partial}_2(\eta_3)$.

- Let P_3 be the path from η_3 to σ_3 via ψ_{32} in Fig. 26. There are three -1's appearing in P_3 and there are three non-critical 2-simplices in P_3. Thus, P_3 contributes $(-1)^{3+3}\sigma_3 = \sigma_3$ to $\tilde{\partial}_2(\eta_3)$.

- Let P_4 be the path from η_3 to σ_4 via ψ_{33} in Fig. 26. There is only one -1 appearing in P_4 and there are three non-critical 2-simplices in P_4. Thus, P_4 contributes $(-1)^{1+3}\sigma_4 = \sigma_4$ to $\tilde{\partial}_2(\eta_3)$.

- Let P_5 be the path from η_3 to σ_4 via ψ_{31} in Fig. 26. There are two -1's appearing in P_5 and there are three non-critical 2-simplices in P_5. Thus, P_5 contributes $(-1)^{2+3}\sigma_4 = -\sigma_4$ to $\tilde{\partial}_2(\eta_3)$.

Therefore, $\tilde{\partial}_2(\eta_3) = -\sigma_1 + \sigma_2 + \sigma_3 + \sigma_4 - \sigma_4 = -\sigma_1 + \sigma_2 + \sigma_3$.

Acknowledgements

References

K. A. Adiprasito, B. Benedetti, and F. H. Lutz. Extremal examples of collapsible complexes and random discrete Morse theory. *Discrete & Computational Geometry*, 57:824–853, 2017. URL `https://doi.org/10.1007/s00454-017-9860-4`.

C. A. Athanasiadis. Decompositions and connectivity of matching and chessboard complexes. *Discrete & Computational Geometry*, 31(3):395–403, 2004. URL `https://doi.org/10.1007/s00454-003-2869-x`.

R. Ayala, D. Fernández-Ternero, and J. A. Vilches. Perfect discrete Morse functions on 2-complexes. *Pattern Recognition Letters*, 33(11):1495–1500, 2012. doi: 10.1016/j.patrec.2011.08.011. URL `https://www.sciencedirect.com/science/article/pii/S0167865511002625`. Computational Topology in Image Context.

B. Benedetti and F. H. Lutz. Random discrete Morse theory and a new library of triangulations. *Experimental Mathematics*, 23(1):66–94, 2014. URL `https://doi.org/10.1080/10586458.2013.865281`.

B. Benedetti, C. Lai, D. Lofano, and F. H. Lutz. Random simple-homotopy theory. *Journal of Applied and Computational Topology*, 2023. URL `https://doi.org/10.1007/s41468-023-00139-4`.

A. Björner, L. Lovász, S. T. Vrećica, and R. T. Živaljević. Chessboard complexes and matching complexes. *Journal of the London Mathematical Society*, 49(1):25–39, 1994. URL `https://doi.org/10.1112/jlms/49.1.25`.

S. Bouc. Homologie de certains ensembles de 2-sous-groupes des groupes symétriques. *Journal of Algebra*, 150(1):158–186, 1992. doi: 10.1016/S0021-8693(05)80054-7. URL `https://www.sciencedirect.com/science/article/pii/S0021869305800547`.

K. S. Brown. Euler characteristics of discrete groups and *G*-spaces. *Inventiones mathematicae*, 27(3): 229–264, 1974. URL `https://doi.org/10.1007/BF01390176`.

K. S. Brown. Euler characteristics of groups: The *p*-fractional part. *Inventiones mathematicae*, 29(1):1–5, 1975. URL `https://doi.org/10.1007/BF01405170`.

M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Mathematics*, 217(1):101–113, 2000. doi: 10.1016/S0012-365X(99)00258-7. URL `https://www.sciencedirect.com/science/article/pii/S0012365X99002587`.

M. M. Cohen. *A Course in Simple-Homotopy Theory*. Springer-Verlag, New York, first edition, 1973.

R. Diestel. *Graph Theory*. Springer-Verlag, New York, fifth edition, 2010.

Ö. Eğecioğlu and T. F. Gonzalez. A computationally intractable problem on simplicial complexes. *Computational Geometry*, 6(2):85–98, 1996. doi: 10.1016/0925-7721(95)00015-1. URL `https://www.sciencedirect.com/science/article/pii/09257721950000151`.

R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134(1):90–145, 1998. doi: 10.1006/aima.1997.1650. URL `https://www.sciencedirect.com/science/article/pii/S0001870897916509`.

R. Forman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire [electronic only]*, 48:B48c–35, 2002.

É. Gallais. Combinatorial realization of the Thom-Smale complex via discrete Morse theory. *Annali della Scuola Normale Superiore di Pisa - Classe di Scienze, Serie 5*, 9(2):229–252, 2010. URL `https://doi.org/10.2422/2036-2145.2010.2.01`.

P. Hersh. On optimizing discrete Morse functions. *Advances in Applied Mathematics*, 35(3):294–322, 2005. doi: 10.1016/j.aam.2005.04.001. URL `https://www.sciencedirect.com/science/article/pii/S0196885805000370`.

J. Jonsson. Five-torsion in the homology of the matching complex on 14 vertices. *Journal of Algebraic Combinatorics*, 29(1):81–90, 2009. URL `https://doi.org/10.1007/s10801-008-0123-6`.

M. Joswig and M. E. Pfetsch. Computing optimal Morse matchings. *SIAM Journal on Discrete Mathematics*, 20(1):11–25, 2006. URL `https://doi.org/10.1137/S0895480104445885`.

K. P. Knudson. *Morse Theory: Smooth and Discrete*. World Scientific, 2015.

D. N. Kozlov. *Organized Collapse: An Introduction to Discrete Morse Theory*. American Mathematical Society, 2020.

T. Lewiner, H. Lopes, and G. Tavares. Toward optimality in discrete Morse theory. *Experimental Mathematics*, 12(3):271–285, 2003a. URL `https://doi.org/10.1080/10586458.2003.10504498`.

T. Lewiner, H. Lopes, and G. Tavares. Optimal discrete Morse functions for 2-manifolds. *Computational Geometry*, 26(3):221–233, 2003b. doi: 10.1016/S0925-7721(03)00014-2. URL `https://www.sciencedirect.com/science/article/pii/S09257721030000142`.

J. R. Munkres. *Elements Of Algebraic Topology*. CRC Press, first edition, 1996.

D. Quillen. Homotopy properties of the poset of nontrivial p-subgroups of a group. *Advances in Mathematics*, 28(2):101–128, 1978. doi: 10.1016/0001-8708(78)90058-0. URL `https://www.sciencedirect.com/science/article/pii/0001870878900580`.

N. A. Scoville. *Discrete Morse Theory*. American Mathematical Society, 2019.

J. Shareshian and M. L. Wachs. Torsion in the matching complex and chessboard complex. *Advances in Mathematics*, 212(2):525–570, 2007. doi: 10.1016/j.aim.2006.10.014. URL `https://www.sciencedirect.com/science/article/pii/S0001870806003719`.

M. L. Wachs. Topology of matching, chessboard, and general bounded degree graph complexes. *Algebra Universalis*, 49(4):345–385, 2003. URL `https://doi.org/10.1007/s00012-003-1825-1`.

J. H. C. Whitehead. Simplicial spaces, nuclei and m-groups. *Proceedings of the London Mathematical Society*, s2-45(1):243–327, 1939. URL `https://doi.org/10.1112/plms/s2-45.1.243`.

E. Zeeman. On the dunce hat. *Topology*, 2(4):341–358, 1963. doi: 10.1016/0040-9383(63)90014-4. URL `https://www.sciencedirect.com/science/article/pii/0040938363900144`.