Subdivisional Spaces and Graph Braid Groups

Byung Hee An, Gabriel C. Drummond-Cole,
and Ben Knudsen

Received: March 16, 2018
Revised: December 6, 2018

Communicated by Max Karoubi

Abstract. We study the problem of computing the homology of the configuration spaces of a finite cell complex X. We proceed by viewing X, together with its subdivisions, as a *subdivisional space*—a kind of diagram object in a category of cell complexes. After developing a version of Morse theory for subdivisional spaces, we decompose X and show that the homology of the configuration spaces of X is computed by the derived tensor product of the Morse complexes of the pieces of the decomposition, an analogue of the monoidal excision property of factorization homology.

Applying this theory to the configuration spaces of a graph, we recover a cellular chain model due to Świątkowski. Our method of deriving this model enhances it with various convenient functorialities, exact sequences, and module structures, which we exploit in numerous computations, old and new.

2010 Mathematics Subject Classification: primary: 20F36, 55R80, 55U05; secondary: 05C10
Keywords and Phrases: configuration spaces, cell complexes, subdivisional spaces, graphs, braid groups

Contents

1 Introduction

Consider the following problem: given a cell structure on a space X, compute the homology of the configuration space

$$B_k(X) \coloneqq \{(x_1, \dots, x_k) \in X^k : x_i \neq x_j \text{ if } i \neq j\}/\Sigma_k,$$

where Σ_k is the symmetric group on $\{1, \dots, k\}$. In this work, we provide new tools to address this problem by combining ideas from homotopy theory and robotics. We apply these tools to study the homology of configuration spaces of graphs.

1.1 Configuration spaces and gluing

First introduced in [FN62], configuration spaces of manifolds have a long and rich history in algebraic topology—for a taste, we direct the reader to [Arn69], [Seg73], [CLM76], [Tot96], and [LS05]. One powerful idea in this area has been the idea that configuration spaces behave predictably under collar gluings; that is, if we decompose a manifold M into two open submanifolds M_0 and M_1 glued along the thickening of an embedded codimension one submanifold N, then the configuration spaces of M are in some sense determined by the configuration spaces of these smaller manifolds. This decomposition technique has borne considerable fruit—see [McD75], [Böd87], [BC88], [BCT89], [FT00], and [Knu17].

One articulation of this gluing property comes from the theory of factorization homology, which provides a quasi-isomorphism

$$C_*(B(M)) \simeq C_*(B(M_0)) \bigotimes^{\mathbb{L}}_{C_*(B(N\times\mathbb{R}))} C_*(B(M_1)),$$

where $B(M) := \coprod_{k\geq 0} B_k(M)$ [AF15]. Unfortunately, these chain complexes are only algebras and modules up to structured homotopy, so this quasi-isomorphism may be difficult to use in computations. To address this issue, we will combine this gluing property with ideas drawn from a different school of thought.

1.2 Configuration spaces and cell structures

Since their introduction in [GK98] and [Ghr02], configuration spaces of graphs have been studied intensively—see [Abr00], [Far03], and [Far05] and the references therein. According to a theorem of [Ghr02], the configuration spaces of a graph Γ are all classifying spaces for their fundamental groups, the so-called *graph braid groups* of Γ.

The key to approaching these spaces is to notice that the cell structure of a graph Γ yields a cellular approximation $B_k^{\square}(\Gamma) \to B_k(\Gamma)$, which becomes a homotopy equivalence after finite subdivision [Abr00]. Thus, one may apply the Morse theory for cell complexes developed in [For98]. This approach has led to many computations—see [FS05], [Far06], [FS08], [FS12], [KKP12], [Sab09], and [KP12].

There are two immediate obstacles to extending the success of cellular methods in the case of a graph to higher dimensional complexes. First, the geometry in higher dimensions may be too difficult to handle directly. To address this problem, we will use the cut-and-paste idea employed in the manifold case, working with locally defined Morse data. Second, the approximation $B_k^{\square}$ may fail to capture the homotopy type of the configuration space in higher dimensions, even after finite subdivision.

1.3 Subdivisional spaces and decomposition

Fortunately, the cellular approximation does improve with subdivision, and, by applying the method of [Abr00], we prove in Theorem 2.8 that, for any suitably *convergent* set $\mathcal{P}$ of subdivisions of X, the map

$$\operatorname*{colim}_{X' \in \mathcal{P}} B_k^{\square}(X') \to B_k(X)$$

is a weak homotopy equivalence. Thus, in the presence of such a convergent *subdivisional structure* $\mathcal{P}$, the study of the configuration spaces of X is equivalent to the study of the diagram of configuration complexes determined by $\mathcal{P}$.

These diagrammatics are formalized within the framework of *subdivisional spaces,* which behave rigidly in some ways and like continuous objects in others. In particular, cellular chains and an abstract form of Morse theory lift effortlessly to the subdivisional context, whereas the *subdivisional configuration space* $B^{\mathrm{SD}}(X)$ captures the homotopy type of $B(X)$. Our first main result is a gluing theorem in this context.

Decomposition Theorem (Theorem 3.20). *Given a decomposition $X \cong X_0 \amalg_{A \times I} X_1$ as subdivisional spaces, together with suitable locally-defined Morse data, the homology of $B(X)$ is computed by the derived tensor product of Morse complexes*

$$I(B^{\mathrm{SD}}(X_0)) \bigotimes^{\mathbb{L}}_{I(B^{\mathrm{SD}}(A \times I))} I(B^{\mathrm{SD}}(X_1)).$$

1.4 Swiatkowski complexes

In the second half of the paper, we apply this theory to the semi-classical case of graphs. What results is a family of chain complexes computing the homology of graph braid groups, which first appeared in the work of Świątkowski [Świ01]. Let Γ be a graph with vertices V and edges E. We set

$$S(\Gamma) \coloneqq \mathbb{Z}[E] \otimes \bigotimes_{v \in V} S(v),$$

where $S(v)$ is the free Abelian group generated by $\{\varnothing, v\} \amalg H(v)$. Here $H(v)$ is the set of half-edges at v. The *Świątkowski complex* of Γ is the (bigraded differential) $\mathbb{Z}[E]$-module $S(\Gamma)$—see §4.2 for details. Our second main result is the following:

Comparison Theorem (Theorem 4.5). *There is a natural isomorphism of bigraded Abelian groups*

$$H_*(B(\Gamma)) \cong H_*(S(\Gamma)).$$

To derive Theorem 4.5 from the decomposition theorem, we fragment Γ completely. We take Γ_0 to be a disjoint union over $v \in V$ of the star graphs $\mathsf{S}_{d(v)}$, where $d(v)$ is the number of half-edges incident on v. We take Γ_1 to be a disjoint union of intervals, one for each edge in E. We obtain Γ by gluing these pieces along $2|E|$ disjoint intervals. We define local Morse data by putting the vertex of each star graph "at the top," so that configurations flow down the legs (see §4.3)—and the resulting Morse complex is a reduced version of the Świątkowski complex for the star graph. The decomposition theorem gives the isomorphism

$$H_*(B(\Gamma)) \cong H_*\left(\bigotimes_{v \in V} \mathbb{Z}[H(v)] \otimes S(v) \bigotimes_{\mathbb{Z}[E]^{\otimes 2}} \mathbb{Z}[E],\, \partial \right),$$

and the righthand side is isomorphic to $S(\Gamma)$ by inspection.

1.5 Homology computations

The Świątkowski complex has many desirable features. It is finite dimensional in each bidegree and finitely generated as a $\mathbb{Z}[E]$-module. It connects configuration spaces of different cardinalities by the action of $\mathbb{Z}[E]$. It depends only on intrinsic graph theoretic data and requires no choice of subdivision. It decomposes geometrically, assigning a short exact sequence to the removal of a vertex. It is functorial for embeddings among graphs, so relations at the level of atomic subgraphs impose global constraints. Some of these properties are evident or implicit already in [Świ01]; several are new. These features amount to a robust computational toolkit, which we exploit extensively in §5 and Appendix C.

1.6 HOW TO READ THIS PAPER

The reader concerned mainly with graph braid groups may wish to start with just enough of §4 to see our conventions on graphs and the definition of the Świątkowski complex before skipping directly to the computations of §5, returning to the theory later. Starting from the beginning is recommended for the reader interested in configuration spaces in general, higher dimensional applications, or variations on the ideas of factorization homology.

1.7 RELATION TO PREVIOUS WORK

This paper grew out of the desire to combine the local-to-global approach to configuration spaces of graphs promised by the stratified factorization homology developed in [AFT17] with the combinatorial character and computational ease of the discrete Morse theoretic model of [FS05], following [Abr00]. Although we do not directly employ the results of any of this work, its ideas permeate the theory developed in §2–4.

The Świątkowksi complex first appeared in [Świ01] (see also [Lüt14]). There Świątkowski constructed a cubical complex lying inside $B_k(\Gamma)$ as a deformation retract to study the fundamental group. The cellular chain complex of this cubical complex is isomorphic to the weight k subcomplex of $S(\Gamma)$. This observation implies a weaker version of Theorem 4.5 which contains more direct geometric content.

A similar edge stabilization mechanism, in a different complex and for trees only, was studied by Ramos [Ram18].

The generators and some of the relations that we describe for the first homology groups appear in [HKRS14]. That work uses these generators to perform many computations motivated by physics, and it contains another alternative proof of the theorem of [KP12] dealt with in our Appendix C.

1.8 FUTURE DIRECTIONS

We defer pursuit of the following ideas to future work.

1. *Edge stabilization.* In the sequel to this paper [ADCK], we show that the $\mathbb{Z}[E]$-action is geometric, arising from a new family of stabilization maps at the level of the configuration spaces themselves, and we carry out a detailed investigation of its properties.

2. *Destabilization.* Dual to the process of adding points is that of splitting configurations apart, which may be phrased as a cocommutative coalgebra structure for which $\mathbb{Z}[E]$ acts by coderivations.

3. *Higher dimensions.* Little research has been done on configuration spaces of higher dimensional cell complexes in general—see [Gal01], [AP17], and [WG] for rare examples. Replicating the computational success of

the Świątkowski complex in higher dimensions amounts to identifying tractable local Morse data.

4. *Cup products.* The diagonal is not a cellular embedding, but it is an embedding of subdivisional spaces, so our methods may shed light on the cohomology rings of configuration spaces. This is already very interesting for graphs—see [Sab09].

5. *Ordered configurations.* Our program translates with minor modifications to the context of ordered configurations, and we expect to recover an enhanced version of the cellular chain complex of the cubical model constructed in [Lüt14].

1.9 Conventions

Graded objects are concentrated in non-negative degrees. This restriction is only used in Proposition A.7. We write $\mathcal{C}h_{\mathbb{Z}}$ for the category of chain complexes. Bigradings of modules are by *degree* and *weight*. The braiding isomorphism for a tensor product of modules has a sign which depends on degree and not on weight: if x and y have degree i and j, the braiding isomorphism takes $x \otimes y$ to $(-1)^{ij} y \otimes x$. We write $[m]$ for the degree shift functor by m and $\{n\}$ for the weight shift functor by n so that the degree i and weight j component of $M[m]\{n\}$ is the degree $i - m$ and weight $j - n$ component of M.
Symmetric monoidal functors are strong monoidal. We use the phrases "(natural) weak equivalence" to refer to a (natural) isomorphism in the relevant homotopy category and the phrases "(natural) quasi-isomorphism" to refer to a (natural) chain map which induces an isomorphism on homology groups.
We write $C^{\mathrm{sing}}(X)$ for the singular chain complex of the topological space X. If X is a CW complex, we denote the cellular chain complex of X by $C(X)$.

1.10 Acknowledgments

The second author thanks Felix Boes for sharing computational tools. The third author thanks Daniel Lütgehetmann for noticing the connection to the work of Świątkowski [Świ01]. The authors thank Tomasz Maciążek for pointing out several related papers.

2 Subdivisional spaces

Following the ideas of [Abr00], we approximate the configuration spaces of a cell complex by cell complexes, and we show, in Theorem 2.8, that the homotopy types of these approximations often "converge" to the homotopy type of the true configuration space under transfinite subdivision. We then introduce the framework of subdivisional spaces, which sets a complex X equipped with a set of subdivisions $\mathcal{P}$ on equal footing with the corresponding collection of configuration complexes $\{B_k^{\square}(X')\}_{X'\in\mathcal{P}}$. We identify a natural theory of homology for these objects, the complex of subdivisional chains, and we show that it is homotopically well-behaved.

2.1 Complexes and subdivision

If X is a CW complex and $c \subseteq X$ is an n-cell, we write ∂c for the image of ∂D^n under a characteristic map for c, and we set $\mathring{c} := c \setminus \partial c$. The n-skeleton of X, which is to say the union of the cells of X of dimension at most n, is denoted $\mathrm{sk}_n(X)$. A cellular map is a map preserving skeleta.
From now on, a *complex* will be a finite CW complex. We choose to restrict our attention to the finite case for the sake of convenience only.

Definition 2.1. Let $f : X \to Y$ be a cellular map between complexes. We say that f is

1. *regular* if f preserves both closed and open cells;
2. an *isomorphism* if f is regular and bijective;
3. an *embedding* if f is regular and injective;
4. a *subdvision* if f is bijective and preserves subcomplexes;
5. a *subdivisional embedding* if f is injective and preserves subcomplexes.

Thus, a subdivisional embedding factors via its image into a subdivision followed by an embedding.
Given a complex X, we write $\mathrm{SD}(X)$ for the category whose objects are subdivisions of X and whose morphisms are commuting triangles of subdivisions. Note that, since a subdivision is in particular a homeomorphism, there can be at most one morphism in $\mathrm{SD}(X)$ with fixed source and target.

Remark 2.2. Some authors consider a subdivision to be the inverse to what we have defined to be a subdivision. We choose this convention because it matches the direction of the functoriality that will arise naturally in the examples of interest. Modulo this issue of direction, our notion of subdivision is equivalent to that of [LW69, Def. II.6.2].

2.2 Configuration complexes

We now introduce the main object of study.

Definition 2.3. Let X be a topological space.

1. The kth ordered *configuration space* of X is

$$\mathrm{Conf}_k(X) = \{(x_1, \ldots, x_k) : x_i \neq x_j \text{ if } i \neq j\}.$$

2. The kth *unordered* configuration space of X is the quotient

$$B_k(X) = \mathrm{Conf}_k(X)/\Sigma_k.$$

Unfortunately, a cell structure on X does not induce an obvious cell structure on $\mathrm{Conf}_k(X)$; however, following [Abr00], there is a cellular approximation.

Definition 2.4. Let X be a complex.

1. The kth ordered *configuration complex* of X is the largest subcomplex $\mathrm{Conf}_k^{\square}(X) \subseteq X^k$ contained in $\mathrm{Conf}_k(X)$.

2. The kth unordered configuration complex of X is the quotient

$$B_k^{\square}(X) = \mathrm{Conf}_k^{\square}(X)/\Sigma_k.$$

In other words, a cell $(c_1, \ldots, c_k)$ of X^k lies in $\mathrm{Conf}_k^{\square}(X)$ if and only if $c_i \cap c_j = \varnothing$ for $i \neq j$.
The quotient $B_k^{\square}(X)$ is again a complex, which we view as an approximation to $B_k(X)$. These approximations enjoy a certain functoriality.

Lemma 2.5. *Let $s : X \to X'$ be a subdivision. The restriction of s^k to $\mathrm{Conf}_k^{\square}(X)$ factors Σ_k-equivariantly through $\mathrm{Conf}_k^{\square}(X')$ as a subdivisional embedding.*

Our next result, which will serve as a replacement for [Abr00, Thm. 2.1], asserts that, after perhaps transfinite subdivision, this combinatorial approximation is faithful. It will be useful to state this result in some generality, and we introduce the following notion in order to do so.

Definition 2.6. Let X be a complex. A *subdivisional structure* on X is a subcategory $\mathcal{P} \subseteq \mathrm{SD}(X)$ that is non-empty, full, and filtered. We say that the subdivisional structure $\mathcal{P}$ is *convergent* if there is a metric on X such that

$$\lim_{X' \in \mathcal{P}} \max_{c \subseteq X'} \mathrm{diam}(c) = 0.$$

Example 2.7. If X is regular, then the barycentric subdivisions of X form a convergent subdivisional structure—see [LW69, Thm. III.1.7].

Theorem 2.8 (Convergence theorem). *For any complex X, any convergent $\mathcal{P} \subseteq \mathrm{SD}(X)$, and any $k \geq 0$, the natural map*

$$\operatorname*{colim}_{X' \in \mathcal{P}} \mathrm{Conf}_k^{\square}(X') \longrightarrow \mathrm{Conf}_k(X)$$

is a homeomorphism and, in particular, a weak homotopy equivalence.

Proof. Convergence implies that any point of $\mathrm{Conf}_k(X)$ lies in $\mathrm{Conf}_k^{\square}(X')$ for some subdivision $X' \in \mathcal{P}$, i.e., the collection $\{\mathrm{Conf}_k^{\square}(X') : X' \in \mathcal{P}\}$ forms a closed cover of the configuration space. $\square$

Since taking homology commutes with the filtered colimit in question, this observation provides a way of understanding the homology of the configuration spaces of X in terms of the homology of its configuration complexes.

Remark 2.9. The content of the convergence theorem is that configuration complexes are useful in the study of configuration spaces whenever X admits a convergent subdivisional structure, and it would be interesting to know for which X this is the case.

Remark 2.10. It is natural to wonder whether finite subdivision suffices to recover the correct homotopy type. We do not address this question here, noting only that the method of proof completely breaks down; indeed, the homotopy groups of $\mathrm{Conf}_2(D^3) \simeq S^2$ are all non-zero above degree one [IMW16].

We close by noting that the configuration complexes interact predictably with disjoint unions.

Lemma 2.11. *There is a natural commuting diagram*

$$\begin{array}{ccc} \coprod_{i+j=k} B_i^{\square}(X) \times B_j^{\square}(Y) & \xrightarrow{\simeq} & B_k^{\square}(X \amalg Y) \\ \downarrow & & \downarrow \\ \coprod_{i+j=k} B_i(X) \times B_j(Y) & \xrightarrow{\simeq} & B_k(X \amalg Y) \end{array}$$

in which the bottom arrow is a homeomorphism and the top a cellular isomorphism.

2.3 Subdivisional spaces

Complexes and subdivisional embeddings form a category which we denote by $\mathcal{C}\mathrm{x}^{\mathrm{SD}}$.

Definition 2.12. A *subdivisional space* is a functor $\mathcal{X} : \mathcal{P} \to \mathcal{C}\mathrm{x}^{\mathrm{SD}}$ with $\mathcal{P}$ a filtered category.

We write $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ for the category whose objects are subdivisional spaces and whose morphisms are given by

$$\mathrm{Hom}_{\mathcal{E}\mathrm{mb}^{\mathrm{SD}}}(\mathcal{X}, \mathcal{Y}) = \lim_{p \in \mathcal{P}} \operatorname*{colim}_{q \in \mathcal{Q}} \mathrm{Hom}_{\mathcal{C}\mathrm{x}^{\mathrm{SD}}}(\mathcal{X}(p), \mathcal{Y}(q)).$$

In other words, $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ is the category of *ind-objects* $\mathrm{Ind}(\mathcal{C}\mathrm{x}^{\mathrm{SD}})$. We shall make very little use of the general theory of ind-objects, but the reader looking for further information on the subject may consult [KS06, Ch. 6].

Remark 2.13. The formula for hom sets in $\mathrm{Ind}(\mathcal{C})$ is derived from the following intuitions:

1. an object of $\mathcal{C}$ should determine an ind-object of $\mathcal{C}$;
2. a general ind-object of $\mathcal{C}$ should be a filtered colimit of objects of $\mathcal{C}$; and
3. objects of $\mathcal{C}$ should be compact as ind-objects.

We say that $\mathcal{X} : \mathcal{P} \to \mathcal{C}\mathrm{x}^{\mathrm{SD}}$ is *indexed* on $\mathcal{P}$. Subdivisional spaces indexed on different categories may be isomorphic.

Example 2.14. A subdivisional structure $\mathcal{P} \subseteq \mathrm{SD}(X)$ determines a subdivisional space $\mathcal{X} : \mathcal{P} \to \mathcal{C}\mathrm{x}^{\mathrm{SD}}$ sending $X \to X'$ to X'.

Most complexes admit many subdivisional structures and many non-isomorphic realizations as subdivisional spaces. Roughly, we imagine that a subdivisional structure $\mathcal{P}$ forces X to be isomorphic to each of the subdivisions contained in $\mathcal{P}$.

Example 2.15. Since the product of filtered categories is filtered, the levelwise Cartesian product and disjoint union of two subdivisional spaces is again a subdivisional space. Note that the former is not the categorical product, nor is the latter the categorical coproduct; indeed, since we work with embeddings, this failure is already present in $\mathcal{C}\mathrm{x}^{\mathrm{SD}}$. More explicitly, there is in general no embedding from a disjoint union extending embeddings on its components because the images may intersect nontrivially. On the other hand, the projections of an embedding into a product are not necessarily themselves embeddings.

Definition 2.16. The *spatial realization* functor $|-|$ is the composite

$$\mathcal{E}\mathrm{mb}^{\mathrm{SD}} = \mathrm{Ind}(\mathcal{C}\mathrm{x}^{\mathrm{SD}}) \to \mathrm{Ind}(\mathcal{T}\mathrm{op}) \xrightarrow{\mathrm{colim}} \mathcal{T}\mathrm{op}.$$

Example 2.17. If $\mathcal{X}$ is any of the subdivisional spaces of Example 2.14, there is a canonical homeomorphism $|\mathcal{X}| \cong X$.

According to Lemma 2.5, configuration complexes are functorial for subdivisional embeddings, so we may make the following definition.

Definition 2.18. Let $\mathcal{X} : \mathcal{P} \to \mathcal{C}\mathrm{x}^{\mathrm{SD}}$ be a subdivisional space. The kth ordered *subdivisional configuration space* of $\mathcal{X}$ is the subdivisional space

$$\mathcal{P} \xrightarrow{\mathcal{X}} \mathcal{C}\mathrm{x}^{\mathrm{SD}} \xrightarrow{\mathrm{Conf}_k^{\square}} \mathcal{C}\mathrm{x}^{\mathrm{SD}}.$$

Similarly, we have the kth unordered subdivisional configuration space $B_k^{\mathrm{SD}}(\mathcal{X})$. We will be particularly interested in this construction when $\mathcal{X}$ comes from a complex X with a convergent subdivisional structure, for in this case Theorem 2.8 gives a weak homotopy equivalence

$$|B_k^{\mathrm{SD}}(\mathcal{X})| \xrightarrow{\sim} B_k(X).$$

It will often be convenient to consider configuration spaces of all finite cardinalities simultaneously, and we set

$$B^{\mathrm{SD}}(\mathcal{X}) \coloneqq \coprod_{k\geq 0} B_k^{\mathrm{SD}}(\mathcal{X})$$

(note that the indicated disjoint union does in fact exist in $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$, since $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ admits filtered colimits). Thus, we have a functor

$$B^{\mathrm{SD}}: \mathcal{E}\mathrm{mb}^{\mathrm{SD}} \to \mathcal{E}\mathrm{mb}^{\mathrm{SD}}.$$

Using Lemma 2.11, we see that B^{SD} naturally carries the structure of a symmetric monoidal functor, where $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ is symmetric monoidal under disjoint union in the domain and Cartesian product in the codomain.

2.4 Subdivisional chains

Since the complex of cellular chains is functorial for cellular maps between complexes, and, in particular, for subdivisional embeddings, we may make the following definition.

Definition 2.19. The functor C^{SD} of *subdivisional chains* is the composite

$$\mathcal{E}\mathrm{mb}^{\mathrm{SD}} = \mathrm{Ind}(\mathcal{C}\mathrm{x}^{\mathrm{SD}}) \xrightarrow{\mathrm{Ind}(C)} \mathrm{Ind}(\mathcal{C}\mathrm{h}_{\mathbb{Z}}) \xrightarrow{\mathrm{colim}} \mathcal{C}\mathrm{h}_{\mathbb{Z}}.$$

Viewing $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ and $\mathcal{C}\mathrm{h}_{\mathbb{Z}}$ as symmetric monoidal under Cartesian product and tensor product, respectively, C^{SD} naturally carries the structure of a symmetric monoidal functor (here we use that cellular chains sends products to tensor products and that the tensor product distributes over colimits).

Remark 2.20. In contrast to the cellular chain complex, $C^{\mathrm{SD}}(\mathcal{X})$ is typically a very large object. For example, if $\mathcal{X}$ is obtained by equipping the interval I with the subdivisional structure $\mathrm{SD}(I)$, then $C^{\mathrm{SD}}(\mathcal{X})$ is uncountably generated.

The fundamental fact about the functor of subdivisional chains is the following.

Proposition 2.21. *There is a natural weak equivalence*

$$C^{\mathrm{SD}}(-) \simeq C^{\mathrm{sing}}(|-|)$$

of functors from $\mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ *to chain complexes.*

In the proof, we use the following intermediary.

Construction 2.22. Let X be a CW complex. We define $S^{\mathrm{cell}}(X)$ to be the simplicial set given in simplicial degree n by the set of *cellular* maps $\sigma: \Delta^n \to X$. This construction is functorial for cellular maps. There is an inclusion $S^{\mathrm{cell}}(X) \to S(X)$, where S is the standard functor of singular simplices, and this map is a weak homotopy equivalence of simplicial sets by the cellular approximation theorem. Since both the induced map on geometric realizations and the composite $|S^{\mathrm{cell}}(X)| \to |S(X)| \to X$ are cellular by construction, we obtain the zig-zag

$$C(X) \xleftarrow{\sim} C(|S^{\mathrm{cell}}(X)|) \xrightarrow{\sim} C(|S(X)|) \cong C^{\mathrm{sing}}(|X|)$$

Proof of Proposition 2.21. Consider the following diagram:

$$\begin{array}{ccccc} \mathrm{Ind}(\mathcal{C}\mathrm{x}^{\mathrm{SD}}) & \longrightarrow & \mathrm{Ind}(\mathcal{T}\mathrm{op}) & \xrightarrow{\mathrm{colim}} & \mathcal{T}\mathrm{op} \\ & \searrow^{\mathrm{Ind}(C(|S^{\mathrm{cell}}|))} \ \searrow_{\mathrm{Ind}(C)} & \downarrow{\scriptstyle \mathrm{Ind}(C^{\mathrm{sing}})} & & \downarrow{\scriptstyle C^{\mathrm{sing}}} \\ & & \mathrm{Ind}(\mathcal{C}\mathrm{h}_{\mathbb{Z}}) & \xrightarrow{\mathrm{colim}} & \mathcal{C}\mathrm{h}_{\mathbb{Z}}. \end{array}$$

The composition along the top and right is $C^{\mathrm{sing}}(|-|)$. The composition along the bottom is C^{SD}. The right square commutes up to natural isomorphism because the colimit is filtered. The triangle and bigon on the left commute up to natural objectwise quasi-isomorphism. Since filtered colimits preserve quasi-isomorphisms, the conclusion follows. □

Thus there is a homotopically well-behaved natural isomorphism

$$H_*(C^{\mathrm{SD}}(\mathcal{X})) \cong H_*(|\mathcal{X}|).$$

To make this statement precise, we use *homotopy colimits* for diagrams of chain complexes—reminders and references are in Appendix A. The following corollary will be a key ingredient in our proof of Theorem 3.20 below.

Corollary 2.23. *Let $\mathcal{X}$ be a subdivisional space and $F : \mathcal{D} \to \mathcal{E}\mathrm{mb}^{\mathrm{SD}}$ a functor equipped with a natural transformation to the constant functor at $\mathcal{X}$. If the induced map $\mathrm{hocolim}_{\mathcal{D}} |F| \to |\mathcal{X}|$ is a weak homotopy equivalence, then the induced map*

$$\operatorname*{hocolim}_{\mathcal{D}} C^{\mathrm{SD}}(F) \to C^{\mathrm{SD}}(\mathcal{X})$$

is a quasi-isomorphism.

Proof. Applying homology to the natural weak equivalence of Proposition 2.21, and using the fact that a levelwise weak homotopy equivalence of functors induces a weak homotopy equivalence on homotopy colimits (Lemma A.4), we obtain the isomorphisms in the following commutative square:

$$\begin{array}{ccc} H_*\left(\mathrm{hocolim}_{\mathcal{D}}\, C^{\mathrm{sing}}(|F|)\right) & \longrightarrow & H_*(C^{\mathrm{sing}}(|\mathcal{X}|)) \\ \downarrow{\scriptstyle \wr} & & \downarrow{\scriptstyle \wr} \\ H_*\left(\mathrm{hocolim}_{\mathcal{D}}\, C^{\mathrm{SD}}(F)\right) & \longrightarrow & H_*(C^{\mathrm{SD}}(\mathcal{X})). \end{array}$$

From our assumption and Proposition A.8, the top map is an isomorphism, and the claim follows. □

3 Decomposition

We prove our first main result, the decomposition theorem, stated below as Theorem 3.20. A careful formulation of this result requires that we supply a certain amount of definitional groundwork, and this task will occupy our attention in §3.1–3.3. The theorem and its proof appear in §3.4.
The proof of the decomposition theorem is premised on various manipulations of homotopy colimits. For the convenience of the reader less familiar with categorical homotopy theory, we have included a brief review of the relevant terminology and results in Appendix A, as well as a number of references.

3.1 Decompositions and gaps

We first make precise the data involved in the type of decomposition that we wish to consider. Before doing so, we remind the reader of two operations on subdivisions. First, by restricting in the source and target, a subdivision yields a subdivision on any subcomplex. Second, the Cartesian product of two subdivisions is a subdivision of the Cartesian product. Both of these constructions respect further subdivision, so a subdivisional structure on a complex yields a subdivisional structure on any subcomplex by restriction, and any subdivisional structures on two complexes yield a subdivisional structure on their product.

Definition 3.1. An (r-fold) *decomposition* of the complex X is the data of

1. a collection of complexes $\{\widetilde{X}_0, \ldots, \widetilde{X}_r, A_1, \ldots, A_r\}$;
2. for each $0 \leq j \leq r$, a pair of embeddings $A_j \to \widetilde{X}_j \leftarrow A_{j+1}$ with disjoint images, where $A_0 = A_{r+1} = \varnothing$ by convention;
3. an isomorphism
$$X \cong \widetilde{X}_0 \coprod_{A_1\times\{0\}} (A_1 \times I) \coprod_{A_1\times\{1\}} \cdots \coprod_{A_r\times\{0\}} (A_r \times I) \coprod_{A_r\times\{1\}} \widetilde{X}_r;$$
4. a subdivisional structure $\mathcal{P}_X \subseteq \mathrm{SD}(X)$ restricting to a product of subdivisional structures $\mathcal{P}_{A_j} \times \mathrm{SD}(I)$ on $A_j \times I$ for each $1 \leq j \leq r$.

We say that the decomposition is *convergent* if $\mathcal{P}_X$ is so.

Given a decomposition, we set
$$X_j := (A_j \times I) \coprod_{A_j\times\{1\}} \widetilde{X}_j \coprod_{A_{j+1}\times\{0\}} (A_{j+1} \times I).$$

These complexes are called the *components* of the decomposition, and the complexes $A_j \times I$ are the *bridges.*

Remark 3.2. Essentially all of our results hold if $\mathrm{SD}(I)$ is replaced with a convergent subdivisional structure on I.

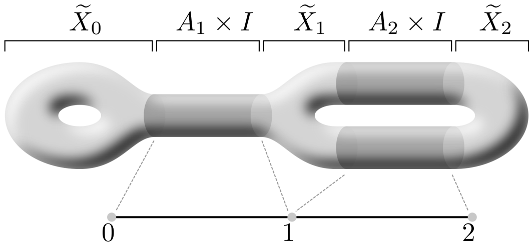

(a) A 2-fold decomposition $\mathcal{E}$ and the map $\pi_{\mathcal{E}}$

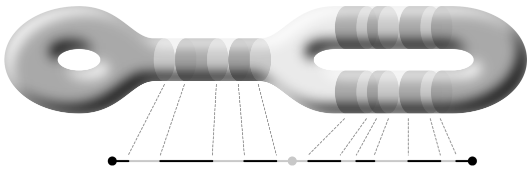

(b) A gap and the complement of its preimage under $\pi_{\mathcal{E}}$

Figure 1: A decomposition and a gap

We typically abbreviate the data of a decomposition to the letter $\mathcal{E}$. Note that, by restriction, a decomposition determines subdivisional structures on all of the complexes involved, and every inclusion between two such lifts to a morphism of subdivisional spaces.

Definition 3.3. Let $\mathcal{E}$ and $\mathcal{F}$ be r-fold decompositions of X and Y, respectively. A *map of decompositions* from $\mathcal{E}$ to $\mathcal{F}$ is a map $f : X \to Y$ of subdivisional spaces whose restrictions fit into a commuting diagram of subdivisional spaces

$$\begin{array}{ccccccccc}
\widetilde{X}_0 & \longleftarrow & A_1 & \longrightarrow & \cdots & \longleftarrow & A_r & \longrightarrow & \widetilde{X}_r \\
\downarrow & & \downarrow & & & & \downarrow & & \downarrow \\
\widetilde{Y}_0 & \longleftarrow & B_1 & \longrightarrow & \cdots & \longleftarrow & B_r & \longrightarrow & \widetilde{Y}_r,
\end{array}$$

and such that $f|_{A_j \times I} = f|_{A_j} \times \mathrm{id}_I$ for each $1 \leq j \leq r$.

A decomposition gives rise to a combinatorial relationship to a certain poset.

Definition 3.4. The category of (r-fold) *gaps* is the partially ordered set $\mathcal{G}_r$ of nonempty open subsets $A \subseteq [0, r]$ such that

1. the complement of A is a (possibly empty) finite union of closed intervals of positive length;

2. if $i \in \{0, \dots, r\}$ lies in the closure of A, then $i \in A$; and

3. for every $1 \leq j \leq r$, $A \cap [j-1, j] \neq \varnothing$.

Construction 3.5. Given a decompositon $\mathcal{E}$ of X, there is a continuous map $\pi_{\mathcal{E}} : X \to [0, r]$ specified by requiring that

1. $\pi_{\mathcal{E}}(\widetilde{X}_j) = \{j\}$, and

2. $\pi_{\mathcal{E}}|_{A_j \times I}$ is the projection onto $I \cong [j-1, j]$.

If $A \subseteq [0, r]$ is a gap, condition (4) of Definition 3.1 provides the inverse image of $[0, r] \setminus A$ with a canonical subdivisional structure, and we obtain this way a functor

$$\gamma_{\mathcal{E}} : \mathcal{G}_r^{op} \to \mathcal{E}\mathrm{mb}^{\mathrm{SD}}.$$

If $f : \mathcal{E} \to \mathcal{F}$ is a map of decompositions, then $f(\pi_{\mathcal{E}}^{-1}([0, r]\setminus A)) \subseteq \pi_{\mathcal{F}}^{-1}([0, r]\setminus A)$ for every $A \in \mathcal{G}_r$, so we may interpret f as a natural transformation from $\gamma_{\mathcal{E}}$ to $\gamma_{\mathcal{F}}$.

Each $\gamma_{\mathcal{E}}(A)$ is a (possibly empty) union of some number of components of the form $(A_j \times [a, 1]) \cup \widetilde{X}_j \cup (A_{j+1} \times [0, b])$ with $0 < a, b < 1$ or $U = A_j \times [c, d]$ with $0 < c < d < 1$. We refer to such a component as a *basic*. We say that the former type of basic is of *component type* and the latter of *bridge type*. The term "basic" is borrowed from [AFT17], whose ideas heavily influence our approach to Theorem 3.20.
Note that X itself typically does not lie in the image of $\gamma_{\mathcal{E}}$, since $\varnothing \notin \mathcal{G}_r$.

3.2 Local invariants

In this paper, our main interest in the decomposition theorem stated below will be as a tool to study configuration spaces, but the proof of the theorem will only make use of a few key features of these spaces.

Definition 3.6. Let $\mathcal{E}$ be a decomposition of X. An $\mathcal{E}$-*local invariant* is a symmetric monoidal functor $F : (\mathcal{E}\mathrm{mb}^{\mathrm{SD}}, \amalg) \to (\mathcal{E}\mathrm{mb}^{\mathrm{SD}}, \times)$ such that the natural map

$$\operatorname*{hocolim}_{\mathcal{G}_r^{op}} |F(\gamma_{\mathcal{E}})| \to |F(X)|$$

is a weak equivalence. A map of local invariants (possibly for different decompositions) is a symmetric monoidal natural transformation.

Remark 3.7. At the cost of greater verbal overhead, it is possible to work with invariants that are only defined locally relative to a given $\mathcal{E}$. All of our results carry over into this more general context.

We now check that this condition is satisfied in the example of greatest interest to us.

Proposition 3.8. *The symmetric monoidal functor B^{SD} is an $\mathcal{E}$-local invariant for any convergent decomposition $\mathcal{E}$.*

Proof. The claim will follow by two-out-of-three after verifying that each of the numbered arrows in the commuting diagram

$$\begin{array}{ccc}
\operatorname*{hocolim}_{A\in\mathcal{G}_r^{op}} B\left(\pi_{\mathcal{E}}^{-1}\left([0,r]\setminus\overline{A}\right)\right) & \xrightarrow{(1)} & B(X) \\
{\scriptstyle (2)}\downarrow & & \Vert \\
\operatorname*{hocolim}_{A\in\mathcal{G}_r^{op}} B\left(|\gamma_{\mathcal{E}}(A)|\right) & \xrightarrow{(3)} & B(X) \\
{\scriptstyle (4)}\uparrow & & \uparrow{\scriptstyle (5)} \\
\operatorname*{hocolim}_{\mathcal{G}_r^{op}} |B^{\mathrm{SD}}(\gamma_{\mathcal{E}})| & \longrightarrow & |B^{\mathrm{SD}}(\mathfrak{X})|,
\end{array}$$

is a weak homotopy equivalence, where $\mathfrak{X}$ is the subdivisional space determined by X and the subdivisional structure $\mathcal{P}_X \subseteq \mathrm{SD}(X)$ of $\mathcal{E}$. The second equivalence follows from the fact that B preserves homotopies through injective maps between injective maps, and the third follows from (1) and (2) by two-out-of-three. Theorem 2.8 gives the fifth and (together with Lemma A.4) the fourth, using our assumption on $\mathcal{E}$ (see the discussion following Definition 2.18). Thus, it remains to verify the first equivalence.

Consider the collection $\mathcal{U} := \{B(\pi_{\mathcal{E}}^{-1}\left([0,r]\setminus\overline{A}\right)) : A \in \mathcal{G}_r\}$, which is an open cover of $B(X)$. In fact, $\mathcal{U}$ is a *complete* cover in the sense of Definition A.5; to see this, we note that for S finite,

$$\begin{aligned}
\bigcap_S B(\pi_{\mathcal{E}}^{-1}\left([0,r]\setminus\overline{A_s}\right)) &= B\left(\bigcap_S \pi_{\mathcal{E}}^{-1}\left([0,r]\setminus\overline{A_s}\right)\right) \\
&= B\left(\pi_{\mathcal{E}}^{-1}\left(\bigcap_S [0,r]\setminus\overline{A_s}\right)\right) \\
&= B\left(\pi_{\mathcal{E}}^{-1}\left([0,r]\setminus\bigcup_S\overline{A_s}\right)\right),
\end{aligned}$$

and that a finite union of closures of gaps is again the closure of a gap. With this observation, the desired equivalence follows from Theorem A.6. □

Remark 3.9. The local invariant B^{SD}, from its definition as a disjoint union over finite cardinalities, is naturally a *graded* subdivisional space. Furthermore, every map induced by an inclusion $U \subseteq V$ automatically preserves this grading, as do the isomorphisms $B^{\mathrm{SD}}(U_1 \amalg U_2) \cong B^{\mathrm{SD}}(U_1) \times B^{\mathrm{SD}}(U_2)$. It is very useful to keep track of this grading in studying configuration spaces (see Theorem 4.5, for example), and we will often do so implicitly.

3.3 Flows and bimodules

The decomposition theorem allows us to study configuration spaces and other local invariants in terms of local information. In order to get a combinatorial handle on the local information, we axiomatize what it means to simplify a local invariant coherently.

DEFINITION 3.10. The *category of abstract flows*, $\mathcal{F}\mathrm{lw}_{\mathbb{Z}}$, has objects (C, π), where C is a chain complex and $\pi : C \xrightarrow{\sim} C$ is an idempotent quasi-isomorphism. A morphism $(C, \pi) \to (C', \pi')$ is a chain map f such that $\pi' \circ f = \pi' \circ f \circ \pi$. Such a map is called *flow compatible*.
We will use two functors $\mathcal{F}\mathrm{lw}_{\mathbb{Z}} \to \mathcal{C}\mathrm{h}_{\mathbb{Z}}$. The *forgetful functor* takes (C, π) to C and is the identity on morphisms. The *Morse complex* I takes (C, π) to $\pi(C)$ and f to $\pi \circ f$. The subcategory of objects of $\mathcal{F}\mathrm{lw}_{\mathbb{Z}}$ with underlying chain complexes flat in each degree inherits a symmetric monoidal structure for which each of these functors is symmetric monoidal.

We think of π as the limit of a flow on C and the elements of the associated Morse complex as the critical points of that flow.

Remark 3.11. A discrete flow in the sense of [For98] gives rise to an abstract flow; indeed, our definitions of abstract flow and flow compatible map were motivated by the desire to work functorially with discrete Morse data.

DEFINITION 3.12. Let $\mathcal{X}$ be a subdivisional space indexed on $\mathcal{P}$.

1. A *subdivisional flow* on $\mathcal{X}$ is the data of the dashed lift in the diagram

$$\begin{array}{ccccc} & & & & \mathcal{F}\mathrm{lw}_{\mathbb{Z}} \\ & & & \nearrow & \downarrow \scriptstyle\text{forget} \\ \mathcal{P} & \xrightarrow{\mathcal{X}} & \mathcal{C}\mathrm{x}^{\mathrm{SD}} & \xrightarrow{C} & \mathcal{C}\mathrm{h}_{\mathbb{Z}}. \end{array}$$

2. A map f of subdivisional spaces equipped with subdivisional flows is *flow compatible* if $C(f)$ lies in the image of the forgetful functor $\mathrm{Ind}(\mathcal{F}\mathrm{lw}_{\mathbb{Z}}) \to \mathrm{Ind}(\mathcal{C}\mathrm{h}_{\mathbb{Z}})$.

DEFINITION 3.13. Let $\mathcal{X} : \mathcal{P} \to \mathcal{C}\mathrm{x}^{\mathrm{SD}}$ be a subdivisional space equipped with a subdivisional flow $g = \{g_p : C(\mathcal{X}(p)) \to C(\mathcal{X}(p))\}_{p \in \mathcal{P}}$. The associated *Morse complex* is the chain complex

$$I(C^{\mathrm{SD}}(\mathcal{X}), g) = \operatorname*{colim}_{\mathcal{P}} I(C(\mathcal{X}(p)), g_p).$$

The Morse complex is functorial for flow compatible maps and comes equipped with a natural quasi-isomorphism

$$C^{\mathrm{SD}}(\mathcal{X}) \xrightarrow{\sim} I(C^{\mathrm{SD}}(\mathcal{X}), g),$$

since a filtered colimit of quasi-isomorphisms is a quasi-isomorphism. In particular, the Morse complex computes the homology of $|\mathcal{X}|$. Typically, when the choice of subdivisional flow g is clear from context, we abbreviate the Morse complex to $I(C^{\mathrm{SD}}(\mathcal{X}))$ or simply $I(\mathcal{X})$.

If $\mathcal{X}$ and $\mathcal{Y}$ are equipped with subdivisional flows, then, since the functor of cellular chains is symmetric monoidal and takes values in levelwise flat chain complexes, the product $\mathcal{X} \times \mathcal{Y}$ inherits a canonical subdivisional flow, which we refer to as the *product flow.* Note that, in this case, there are canonical isomorphisms $I(\mathcal{X} \times \mathcal{Y}) \cong I(\mathcal{X}) \otimes I(\mathcal{Y})$ satisfying obvious associativity and commutativity relations.

DEFINITION 3.14. Let $\mathcal{E}$ be a decomposition of X and F an $\mathcal{E}$-local invariant. A *local flow* on F is the data of a subdivisional flow on $F(\gamma_{\mathcal{E}}(A))$ for each $A \in \mathcal{G}_r$, subject to the following conditions:

1. for $A \subseteq B$, the induced map $F(B) \to F(A)$ is flow compatible; and
2. the isomorphism $F(\gamma_{\mathcal{E}}(A)) \times F(\gamma_{\mathcal{E}}(B)) \cong F(\gamma_{\mathcal{E}}(A) \amalg \gamma_{\mathcal{E}}(B))$ and its inverse are each flow compatible, where the lefthand side carries the product flow, whenever $\gamma_{\mathcal{E}}(A)$ and $\gamma_{\mathcal{E}}(B)$ are disjoint in X.

A map between local invariants equipped with local flows is *flow compatible* if each of its components is so.

A local flow in particular determines the dashed lift in the diagram

$$\begin{array}{ccccccc} & & & & & & \mathrm{Ind}(\mathcal{F}\mathrm{lw}_{\mathbb{Z}}) \\ & & & & & \nearrow & \downarrow \mathrm{colim} \\ \mathcal{G}_r^{op} & \longrightarrow & \mathcal{E}\mathrm{mb}^{\mathrm{SD}} & \xrightarrow{F} & \mathcal{E}\mathrm{mb}^{\mathrm{SD}} & \xrightarrow{C^{\mathrm{SD}}} & \mathcal{C}\mathrm{h}_{\mathbb{Z}}. \end{array}$$

Since the homotopy colimit of the bottom composite computes the homology of $|F(X)|$ by Corollary 2.23 and Definition 3.6, we can hope to understand this homology by means of the Morse complexes associated to these subdivisional flows. We now introduce a further condition guaranteeing that the Morse theory is sufficiently rigid.

DEFINITION 3.15. Let $\mathcal{E}$ be a decomposition of X and F an $\mathcal{E}$-local invariant. A local flow on F is *isotopy invariant* if every inclusion between two basics of bridge type or two basics of component type induces an isomorphism on Morse complexes.

Note that a local flow does not assign a subdivisional flow to $F(X)$. We also refrain from assigning flows to the components or the bridges; instead, we set $I(F(A_j \times I)) := \mathrm{colim}\, I(F(U))$, where the colimit is taken over all basics contained in $A_j \times I$, and similarly for $I(F(X_j))$. Isotopy invariance guarantees each of these "Morse complexes" is canonically isomorphic to the Morse complex of any corresponding basic.

In a more homotopical context, it would be sensible to require a weaker condition. The strictness of isotopy invariance is motivated by the computational nature of our goals, and it has the following important consequence, which is drawn from the theory of factorization algebras—see [Gin13] and the references therein.

Construction 3.16. Let $\mathcal{E}$ be a decomposition of X and F an $\mathcal{E}$-local invariant equipped with an isotopy invariant local flow. For each $1 \leq j \leq r$, the Morse complex $I(F(A_j \times I))$ of the bridge carries a natural associative algebra structure, for which the Morse complexes $I(F(X_{j-1}))$ and $I(F(X_j))$ of the relevant components are right and left modules, respectively.
We indicate how the algebra structure arises (the module structures are similar). Set $R_j := I(F(A_j \times I))$; then the unit map $\eta_j : \mathbb{Z} \to R_j$ is induced by the inclusion $\varnothing \to A_j \times I$. As for the multiplication, we take any two disjoint subbasics U_1 and U_2 of bridge type contained in $A_j \times I$ and define μ_j to be the composite

$$\begin{aligned} R_j \otimes R_j &\xrightarrow{\simeq} I(F(U_1)) \otimes I(F(U_2)) && (3.15) \\ &\xrightarrow{\simeq} I(F(U_1) \times F(U_2)) \\ &\xrightarrow{\simeq} I(F(U_1 \amalg U_2)) && (3.14(2)) \\ &\longrightarrow I(F(A_j \times I)) = R_j, && (3.14(1)) \end{aligned}$$

where in the second line we have used that the Morse complex of a product subdivisional flow is the tensor product of the Morse complexes.
Any two choices for U_1 and U_2 may be connected by a zig-zag of inclusions of disjoint pairs of basics of bridge type, so, by naturality of the maps in question, μ_j is independent of this choice. Tracing through the construction shows that η_j is a unit for μ_j. For associativity, we consider a configuration of five bridge type basics with containments as in Figure 2. We use U_1 and U_{23} to define the outer multiplication and U_2 and U_3 for the inner multiplication in the expression $\mu_j \circ (\mathrm{id} \otimes \mu_j)$, and we use U_{12} and U_3 to define the outer multiplication and U_1 and U_2 to define the inner multiplication in the expression $\mu_j \circ (\mu_j \otimes \mathrm{id})$. Then both expressions are given in terms of maps $I(F(U_1)) \otimes I(F(U_2)) \otimes I(F(U_3)) \to R_j$, which coincide because of the associativity of the structure morphisms for local invariants.

It will be convenient to have terminology for this emergent algebraic structure.

Definition 3.17. An *r-fold bimodule* is

1. a collection $(R_1, \ldots, R_r, M_0, \ldots, M_r)$ of chain complexes,
2. the structure of a differential graded unital associative algebra on R_j for each $1 \leq j \leq r$, and
3. the structure of an (R_j, R_{j+1})-bimodule on M_j for each $0 \leq j \leq r$, where $R_0 = R_{r+1} = \mathbb{Z}$ by convention.

A map from the r-fold bimodule $(R_1, \ldots, R_r, M_0, \ldots, M_r)$ to the r-fold bimodule $(R'_1, \ldots, R'_r, M'_0, \ldots, M'_r)$ consists of a collection of maps of algebras $R_j \to R'_j$ together with a map of (R_j, R_{j+1})-bimodules from M_j to M'_j for each j, where M'_j carries the bimodule structure induced by restriction along the maps $R_j \to R'_j$ and $R_{j+1} \to R'_{j+1}$.

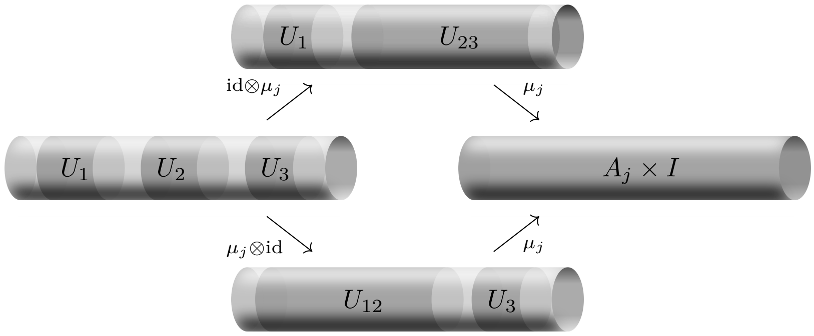

Figure 2: Basics showing the associativity of the product.

It is clear from Construction 3.16 and the definition of a local flow on a local invariant that a flow compatible map between local invariants intertwines the various algebra and module structures.

We summarize the discussion so far in categorical terms. There is a category $\mathcal{I}\mathrm{nvt}_r^{\mathrm{iso}}$ whose objects are triples $(X, \mathcal{E}, F)$ consisting of a complex X, an r-fold decomposition $\mathcal{E}$, and an $\mathcal{E}$-local invariant F equipped with an isotopy invariant local flow (suppressed in the notation), and whose morphisms are pairs of a map of decompositions and a flow compatible map of local invariants. There is a second category $\mathcal{B}\mathrm{imod}_r$ whose objects are r-fold bimodules and whose morphisms are maps of such. What we have constructed so far is a canonical lift of the Morse complex to a functor

$$I : \mathcal{I}\mathrm{nvt}_r^{\mathrm{iso}} \to \mathcal{B}\mathrm{imod}_r.$$

Remark 3.18. In fact, the r-fold bimodule arising from an isotopy invariant local flow is always a *pointed* r-fold bimodule, which is to say that each M_j is equipped with a distinguished map $\mathbb{Z} \to M_j$ arising from the inclusion $\varnothing \to X_j$. Except in defining the natural transformation of Construction 3.22, this extra structure will play little role in what follows.

3.4 The decomposition theorem

Our goal is to use the algebraic structures of Construction 3.16 to express the global value of a local invariant with an isotopy invariant local flow in terms of its values on the pieces of the decomposition. In order to make a precise statement, we first need to spell out how the various pieces of an r-fold bimodule may be assembled to give a corresponding "global value."

Definition 3.19. Let $(R_1, \ldots, R_r, M_0, \ldots, M_r)$ be an r-fold bimodule.

1. The *simplicial bar construction* on these data is the r-fold simplicial chain complex given in degree $(n_1, \ldots, n_r)$ by

$$\mathrm{Bar}_\Delta(M_0, R_1, \ldots, R_r, M_r)_{(n_1,\ldots,n_r)} = M_0 \otimes R_1^{\otimes n_1} \otimes \cdots \otimes R_r^{\otimes n_r} \otimes M_r.$$

The face maps are defined by the respective module action and algebra multiplication maps, and the degeneracies are defined by the respective units.

2. The *bar construction* or *bar complex* on these data is the total chain complex of the multicomplex obtained from the simplicial bar construction by taking the respective alternating sums of the face maps in each simplicial direction.

The bar complex of the r-fold bimodule is a functor in a straightforward manner and computes the homology of the derived tensor product

$$M_0 \otimes^{\mathbb{L}}_{R_1} \cdots \otimes^{\mathbb{L}}_{R_r} M_r.$$

Indeed, depending on one's point of view, this prescription may even be taken as the definition of the derived tensor product. For details on these matters, the reader may consult [Smi67], for example.
We are now equipped to state our main result.

THEOREM 3.20 (Decomposition theorem). *There is a natural weak equivalence connecting the two composites in the diagram*

$$\begin{array}{ccc} \mathcal{I}\mathrm{nvt}_r^{\mathrm{iso}} & \xrightarrow{(X,\mathcal{E},F)\mapsto F(X)} & \mathcal{E}\mathrm{mb}^{\mathrm{SD}} \\ {\scriptstyle I}\downarrow & & \downarrow{\scriptstyle C^{\mathrm{SD}}} \\ \mathcal{B}\mathrm{imod}_r & \xrightarrow{\mathrm{Bar}} & \mathcal{C}\mathrm{h}_{\mathbb{Z}}. \end{array}$$

In particular, there is a weak equivalence

$$C^{\mathrm{SD}}(F(X)) \simeq I(F(X_0)) \bigotimes^{\mathbb{L}}_{I(F(A_1\times I))} \cdots \bigotimes^{\mathbb{L}}_{I(F(A_r\times I))} I(F(X_r)).$$

We turn now to the proof of Theorem 3.20. For the sake of brevity, when considering the r-fold bimodule arising from an $\mathcal{E}$-local invariant F, we use the somewhat abusive notation $\mathrm{Bar}(I(\mathcal{E}))$ for the corresponding bar complex. Since $\mathrm{Bar}(I(\mathcal{E}))$ arises from the multi-simplicial object $\mathrm{Bar}_\Delta(I(\mathcal{E}))$, and since $C^{\mathrm{SD}}(F(X))$ may be recovered as a homotopy colimit over $\mathcal{G}_r^{op}$, our strategy in relating these two objects will be to relate the categories Δ^r and $\mathcal{G}_r$.

DEFINITION 3.21. Let $A \subseteq [0, r]$ be a gap. The jth *trace* of A is the ordered set

$$\tau_j(A) = \pi_0\left(A \cap [j-1, j]\right),$$

with the ordering induced by the standard orientation of $\mathbb{R}$.

It follows from the definitions that the set $\tau_j(U)$ is always non-empty, so the various traces extend to a functor

$$\tau : \mathcal{G}_r \to \Delta^r.$$

Using the trace, we may relate the Morse complex to the bar construction.

CONSTRUCTION 3.22. We define a natural transformation

$$\psi : I(F(\gamma_{\mathcal{E}})) \to \mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op}$$

of functors from $\mathcal{G}_r^{op}$ to chain complexes. For a gap A, we have

$$\begin{aligned}\mathrm{Bar}_\Delta(I(\mathcal{E}))(\tau^{op}(A)) \cong I(F(X_0)) \otimes I(F(A_1 \times I))^{\otimes|\tau_1(A)|-1} \otimes \cdots \\ \cdots \otimes I(F(A_r \times I))^{\otimes|\tau_r(A)|-1} \otimes I(F(X_r))\end{aligned}$$

and we define the component ψ_A by expressing $\gamma_{\mathcal{E}}(A)$ as a disjoint union of basics and tensoring together the maps induced by the inclusions of each basic into the corresponding component or bridge, or of the empty set into the relevant component—see Remark 3.18. Naturality follows from flow compatibility of the structure maps involved, together with the fact that these Morse complexes arise from product flows, both of which are guaranteed by Definition 3.14.

As a matter of terminology, we say that an r-fold gap $A \in \mathcal{G}_r$ is *separated* if $A \cap \{0, \dots, r\} = \varnothing$. Write $\widetilde{\mathcal{G}}_r \subseteq \mathcal{G}_r$ for the full subcategory of separated gaps.

LEMMA 3.23. *Let $\mathcal{E}$ be a decomposition of X and F an $\mathcal{E}$-local invariant with an isotopy-invariant local flow. The canonical map*

$$\operatorname*{hocolim}_{\mathcal{G}_r^{op}} I(F(\gamma_{\mathcal{E}})) \xrightarrow{\sim} \operatorname*{hocolim}_{\mathcal{G}_r^{op}} \mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op}$$

induced by ψ is a quasi-isomorphism.

Proof. A gap A is separated if and only if $\gamma_{\mathcal{E}}(A)$ intersects each $\widetilde{X}_j$ nonvacuously, so, by isotopy invariance, ψ_A is a quasi-isomorphism for separated A. Thus, by Lemma A.4 and Proposition A.12, it suffices to note that the inclusion $\widetilde{\mathcal{G}}_r \subseteq \mathcal{G}_r$ is homotopy initial (so that the inclusion of opposite categories is homotopy final). Indeed, all of the overcategories in question are cofiltered and hence contractible. □

Theorem 3.20 relies on the following fact about the functor τ.

LEMMA 3.24. *For any object $S \in \Delta^r$, the inclusion $\iota : \tau^{-1}(S) \to (S \downarrow \tau)$ is homotopy initial.*

In one form or another, this fact is certainly well-known to experts. In the name of a self-contained narrative, we nevertheless include a proof, which is deferred to §3.5 below. For now, we draw the following consequence (see Appendix A for notation).

COROLLARY 3.25. *Let* $V : (\Delta^{op})^r \to \mathcal{C}h_{\mathbb{Z}}$ *be a multi-simplicial chain complex. There is a natural weak equivalence* $\mathrm{holan}_{\tau^{op}}(V \circ \tau^{op}) \simeq V$.

Proof. We have that

$$\begin{aligned}
\mathrm{hoLan}_{\tau^{op}}(V \circ \tau^{op})(S) &= \operatorname*{hocolim}_{(\tau^{op}\downarrow S)} (V \circ \tau^{op} \circ \mathrm{forget}) && \text{(A.9)}\\
&\simeq \operatorname*{hocolim}_{(\tau^{op})^{-1}(S)} (V \circ \tau \circ \mathrm{forget} \circ \iota) && \text{(3.24, A.12)}\\
&= \operatorname*{hocolim}_{(\tau^{op})^{-1}(S)} (V \circ \underline{S}) \\
&\simeq V(S) && \text{(A.13)}
\end{aligned}$$

where in the last step we have used that the category $\tau^{-1}(S)$ is contractible. To see why this is so, we note that $(S\downarrow\tau)$ is contractible, having a final object (since $\mathcal{G}_r$ has the final object $[0, r]$), and invoke Corollary A.14 and Lemma 3.24 a second time. □

Proof of Theorem 3.20. We have the following column of quasi-isomorphisms:

$$\begin{aligned}
C^{\mathrm{SD}}(F(X)) &\simeq \operatorname*{hocolim}_{\mathcal{G}_r^{op}} C^{\mathrm{SD}}(F(\gamma_{\mathcal{E}})) && \text{(2.23)}\\
&\simeq \operatorname*{hocolim}_{\mathcal{G}_r^{op}} I(F(\gamma_{\mathcal{E}})) && \text{(A.4)}\\
&\simeq \operatorname*{hocolim}_{\mathcal{G}_r^{op}} \tau^*\mathrm{Bar}_\Delta(I(\mathcal{E})) && \text{(3.23, A.4)}\\
&\cong \mathrm{hoLan}_*(\mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op})(*) && \text{(A.10)}\\
&\simeq \mathrm{hoLan}_*\, \mathrm{hoLan}_\tau(\mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op})(*) \\
&\cong \operatorname*{hocolim}_{(\Delta^{op})^r} \mathrm{hoLan}_\tau(\mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op}) && \text{(A.10)}\\
&\simeq \operatorname*{hocolim}_{(\Delta^{op})^r} \mathrm{Bar}_\Delta(I(\mathcal{E})) && \text{(3.25, A.4)}\\
&\simeq \mathrm{Bar}(I(\mathcal{E})) && \text{(A.7)}.
\end{aligned}$$

The unmarked quasi-isomorphism is formal, following from the fact that left Kan extensions compose. Naturality follows from flow compatibility and inspection of Construction 3.22. □

3.5 PROOF OF LEMMA 3.24

The fundamental observation in the proof is the following. Let $\widetilde{\mathcal{G}}_{1,k} \subseteq \widetilde{\mathcal{G}}_1$ denote the (non-full) subcategory with objects the separated gaps with exactly k components and morphisms the π_0-bijective inclusions.

LEMMA 3.26. *For every* $k > 0$, *the category* $\widetilde{\mathcal{G}}_{1,k}$ *is contractible.*

Proof. We define a functor $\chi : \widetilde{\mathcal{G}}_1 \to \mathcal{T}op$ by letting $\chi(A) \subseteq B_k(A)$ be the subspace of configurations that intersect each connected component of A nontrivially. The proof will be complete upon establishing the chain of weak homotopy equivalences

$$|N\widetilde{\mathcal{G}}_{1,k}| \simeq \operatorname*{hocolim}_{\widetilde{\mathcal{G}}_{1,k}} \underline{\mathrm{pt}} \simeq \operatorname*{hocolim}_{\widetilde{\mathcal{G}}_{1,k}} \chi \simeq \mathrm{pt},$$

where $\underline{\mathrm{pt}}$ is the constant functor with value a singleton. The first equivalence is immediate from Definition A.3, and the second follows from Lemma A.4 and the fact that $\chi(A)$ is homeomorphic to the product of the connected components of A and hence contractible. To establish the third equivalence, we note that the collection $\{\chi(A) : A \in \widetilde{\mathcal{G}}_{1,k})\}$ of open subsets of $B_k((0,1))$ is a basis for its topology (and thus a complete cover). Since this configuration space is contractible, the desired equivalence now follows from Theorem A.6. □

Now, a separated r-fold gap is nothing more or less than an r-tuple of separated 1-fold gaps. In other words, there is a commuting diagram of functors

$$\begin{array}{ccccc}
\widetilde{\mathcal{G}}_1^r & \xleftarrow{\sim} & \widetilde{\mathcal{G}}_r & \longrightarrow & \mathcal{G}_r \\
\downarrow & & & & \downarrow \tau \\
\mathcal{G}_1^r & & \xrightarrow{\quad \tau^r \quad} & & \Delta^r.
\end{array}$$

Writing $\widetilde{\tau} : \widetilde{\mathcal{G}}_r \to \Delta^r$ for either composite and fixing an object $S = (S_1, \ldots, S_r) \in \Delta^r$, we have another commuting diagram of functors of the form

$$\begin{array}{ccc}
\tau^{-1}(S) & \xrightarrow{\quad \iota \quad} & (S \downarrow \tau) \\
\uparrow & & \uparrow \\
\widetilde{\tau}^{-1}(S) & \xrightarrow{\quad \widetilde{\iota} \quad} & (S \downarrow \widetilde{\tau}) \\
\wr\downarrow & & \downarrow\wr \\
\prod_{j=1}^r \widetilde{\tau}^{-1}(S_j) & \xrightarrow{\widetilde{\iota}^r} & \prod_{j=1}^r (S_j \downarrow \widetilde{\tau}).
\end{array}$$

The strategy will be to understand ι by understanding each of the other arrows in the diagram (a direct comparison is possible but somewhat more involved), beginning with the upper vertical arrows.

Lemma 3.27. *The inclusions of $\widetilde{\tau}^{-1}(S)$ and $(S \downarrow \widetilde{\tau})$ into $\tau^{-1}(S)$ and $(S \downarrow \tau)$, respectively, are each homotopy initial.*

Proof. We give the proof for the former inclusion only, the latter differing only in requiring more notation. Fixing $A \in \tau^{-1}(S)$, we must check the contractibility of the category of r-fold gaps B such that

1. $B \subseteq A$,

2. $\tau(B \subseteq A) = \mathrm{id}_S$, and

3. B is separated.

First, we note that this category is non-empty; indeed, we may obtain such a B from A by removing a sufficiently small neighborhood of each $j \in \{0, \ldots, r\}$ from A. Moreover, any B satisfying these three conditions is contained in one of this form. Since this subcollection is clearly filtered, the claim follows by Example A.2. □

Next, we consider the middle horizontal arrow.

Lemma 3.28. *The functor $\widetilde{\iota}$ is homotopy initial.*

Proof. The property of being homotopy initial is preserved by products and equivalences of categories, so it suffices to consider the case $r = 1$. We establish some notation.

1. An object of $(S{\downarrow}\widetilde{\tau})$ is a pair (A, f), where $f : S \to T$ is a map of ordered sets and A is a union of open subintervals of $(0, 1)$ whose components have disjoint closures and such that $\pi_0(A) = T$ as ordered sets.

2. A morphism is an inclusion making the evident triangle in Δ commute.

3. The functor $\widetilde{\iota}$ is defined by sending $A \in \tau^{-1}(S)$ to the pair (A, id_S).

We wish to prove the contractibility, for each (A, f), of the category $(\widetilde{\iota}{\downarrow}(A, f))$, which is nothing other than the category of gaps $B \subseteq A$ with $\pi_0(B) = S$ and $\pi_0(B \subseteq A) = f$. By inspection, we have the isomorphism

$$(\widetilde{\iota}{\downarrow}(A, f)) \cong \prod_{t \in \mathrm{im}(f)} \widetilde{\mathcal{G}}_{1,|f^{-1}(t)|},$$

so contractibility follows from Lemma 3.26. □

Proof of Lemma 3.24. It is a fact that homotopy initial functors satisfy a partial two-out-of-three property; that is, if T_1 is homotopy initial, then $T_2 \circ T_1$ is homotopy initial if and only if T_2 is so (see [Lur09, Prop. 4.1.1.3(2)], for example). Applying this fact to the composite $\widetilde{\tau}^{-1}(S) \to \tau^{-1}(S) \to (S \downarrow \tau)$ and invoking Lemma 3.27 reduces the lemma to verifying that the composite is homotopy initial. This composite coincides with the composite $\widetilde{\tau}^{-1}(S) \to (S{\downarrow}\widetilde{\tau}) \to (S{\downarrow}\tau)$, so this claim follows from Lemmas 3.27 and 3.28, since homotopy initial functors compose. □

4 Application to graphs

We use the theory developed above to study the configuration spaces of a graph. We establish conventions, introduce the Świątkowski complex, and use the technology developed in the previous section to prove Theorem 4.5, which asserts that this complex computes the homology of the configuration spaces of a graph functorially.

4.1 Conventions on graphs

A *graph* is a finite 1-dimensional CW complex Γ. Its 0-cells and 1-cells are its *vertices* and *edges* $V(\Gamma)$ and $E(\Gamma)$, or simply V and E. The vertices of an edge $V(e)$ are the vertices contained in the closure of that edge in Γ. We write $E(v)$ for the set of edges incident to the vertex v. A *half-edge* is an end of an edge. The set of half-edges of Γ is $H(\Gamma)$ or simply H. We write $H(v)$ for the set of half-edges incident to v and $H(e)$ for the set of half-edges contained in e. For h in H, we write $v(h)$ and $e(h)$ for the corresponding vertex and edge.
Any sufficiently small neighborhood of a vertex v is homeomorphic to a cone on finitely many points; this finite number is the *valence* of v, denoted $d(v)$. The vertex v is *isolated* if $d(v) = 0$. An edge with a 1-valent vertex is a *tail*. A *self-loop* at a vertex is an edge whose entire boundary is attached at that vertex. A graph is *simple* if it has no self-loops and no pair of edges with the same vertices.

Example 4.1. The cone on $\{1, \ldots, n\}$ is a graph S_n with $n+1$ vertices. These graphs are called *star graphs* and the cone point the *star vertex*.
The interval I is a graph with two vertices 0 and 1 and one edge between them. It is isomorphic to S_1 and homeomorphic to S_2 but it will be convenient to have alternate notation for it.

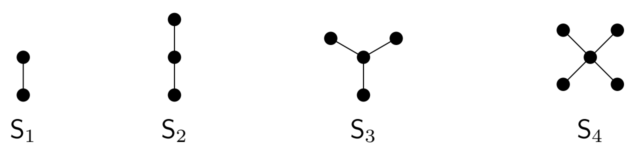

Figure 3: Star graphs

Definition 4.2. Let $f : \Gamma_1 \to \Gamma_2$ be a continuous map between graphs. We say that f is a *graph morphism* if

1. the inverse image $f^{-1}(V(\Gamma_2))$ is contained in $V(\Gamma_1)$ and
2. the map f is injective.

We call a graph morphism a *smoothing* if it is a homeomorphism and a *graph embedding* if it preserves vertices. A graph morphism can be factored into a graph embedding followed by a smoothing. The composite of graph morphisms is a graph morphism, and we obtain in this way a category $\mathcal{G}\mathrm{ph}$. Although the objects of $\mathcal{G}\mathrm{ph}$ are simply finite 1-dimensional CW complexes, not all morphisms are cellular. A *subgraph* is the image of a graph embedding. A graph morphism $f : \Gamma_1 \to \Gamma_2$ induces a map $E(f) : E(\Gamma_1) \to E(\Gamma_2)$, a partially defined map $V(f) : f^{-1}(V(\Gamma_2)) \to V(\Gamma_2)$, and a map $H(v)(f) : H(v) \to H(f(v))$ for each $v \in f^{-1}(V(\Gamma_2))$.

Figure 4: There is a graph morphism (in fact a smoothing) from left to right but not from right to left.

Since graph morphisms are injective, they induce maps at the level of configuration spaces. Thus, it is natural to view $H_*(B(-))$ as a functor from the category $\mathcal{G}\mathrm{ph}$ to bigraded Abelian groups (with a weight grading for cardinality).

4.2 The Swiatkowski complex

We now introduce our main tool in the study of $H_*(B(\Gamma))$.

Construction 4.3 (Świątkowski complex). Let Γ be a graph. For each vertex $v \in V$, we set $S(v) = \mathbb{Z}\langle \varnothing, v, h \in H(v)\rangle$ and regard this Abelian group as bigraded with $|\varnothing| = (0,0)$, $|v| = (0,1)$, and $|h| = (1,1)$.
The *Świątkowski complex* of Γ is the differential bigraded $\mathbb{Z}[E]$-module

$$S(\Gamma) = \mathbb{Z}[E] \otimes \bigotimes_{v \in V} S(v),$$

where $|e| = (0,1)$ for $e \in E$, and differential determined by setting

$$\partial(h) = e(h) - v(h).$$

Since ∂ is $\mathbb{Z}[E]$-linear, the module structure descends to homology.
A graph morphism $f : \Gamma_1 \to \Gamma_2$ determines a map $S(f) : S(\Gamma_1) \to S(\Gamma_2)$. This takes edges to their images under f. If $f(v)$ is a vertex of Γ_2, then the induced map takes $S(v)$ to $S(f(v))$ using f. If $f(v)$ is in the edge e in Γ_2, the map factors through $S(v) \to \mathbb{Z}[e]$ where $\varnothing$ goes to 1, v to e, and $h \in H(v)$ to 0. By inspection, $S(f)$ respects the bigrading, differential, and module structures.

Remark 4.4. The generators of $S(\Gamma)$ describe "states" in the configuration spaces of Γ. The module $S(v)$ records the local states allowed at the vertex v, with $\varnothing$ corresponding to the absence of a particle at v, the element v to a stationary particle at v, and the element $h \in H(v)$ to a path in which a particle moves infinitesimally along the edge containing h. A general state is obtained by prescribing a local state at each vertex and a number of particles on each edge, and the differential is the cellular differential taking a path to its endpoints.

We view S as a functor from $\mathcal{G}\mathrm{ph}$ to the category of bigraded chain complexes and the action of a (weight-graded) ring. A morphism is a weight-graded morphism of rings and a compatible morphism of differential bigraded modules. We denote the degree i and weight k component of $S(\Gamma)$ by $S_i(\Gamma)_k$. Our main result concerning the Świątkowski complex is the following:

Theorem 4.5 (Comparison theorem). *There is an isomorphism*

$$H_*(B(\Gamma)) \cong H_*(S(\Gamma))$$

of functors from $\mathcal{G}\mathrm{ph}$ *to bigraded Abelian groups.*

Remark 4.6. The weight k subcomplex $S(\Gamma)_k$ is isomorphic to the cellular chains of the cubical complex exhibited in [Świ01], if every vertex is at least trivalent. This implies Theorem 4.5 at the level of objects in this special case. The full proof (in particular, functoriality) constitutes the content of this section.
Here we present the $\mathbb{Z}[E]$-module structure on the Świątkowski complex algebraically. In subsequent work [ADCK], we show that this structure arises from an E-indexed family of maps of topological spaces $B_k(\Gamma) \to B_{k+1}(\Gamma)$ that increase the number of points on an edge. Such stabilization maps were known to exist for tails [AP17] and for trees at the level of Morse complexes [Ram18], but stabilization at arbitrary edges is new, and the sequel is devoted to the study of its properties.

It is often useful to consider a smaller variation on the Świątkowski complex.

Definition 4.7. Let Γ be a graph and U a subset of $V(\Gamma)$. For each $v \in U$, let $\widetilde{S}(v) \subseteq S(v)$ be the subspace spanned by $\varnothing$ and the differences $h_{ij} := h_i - h_j$ of half-edges. The *reduced Świątkowski complex* (relative to U), is

$$S^U(\Gamma) := \mathbb{Z}[E] \otimes \bigotimes_{v \in V \setminus U} S(v) \otimes \bigotimes_{v \in U} \widetilde{S}(v),$$

considered as a subcomplex and submodule of $S(\Gamma)$. To be explicit, the differential is determined by $\partial(h_{ij}) = e(h_i) - e(h_j)$.

Notation 4.8. When $U = V$ is the full set of vertices, we write $\widetilde{S}(\Gamma) := S^V(\Gamma)$. When U is the set of 1-valent vertices, we write $S^\partial(\Gamma) := S^U(\Gamma)$. Both $\widetilde{S}(-)$ and $S^\partial(-)$ are functorial for graph morphisms.
When Γ is a disjoint union of star graphs and intervals, we write $S^\curlywedge(\Gamma)$ (an abuse of notation) for the reduced Świątkowski complex of S_n relative to the set of non-star vertices. Since S_1 and the interval are isomorphic, this requires a specification of which such components are star graphs (and which vertex is the star vertex). The construction $S^\curlywedge(-)$ is functorial only for graph morphisms $f : \Gamma_1 \to \Gamma_2$ such that $f^{-1}(v)$ is a star vertex whenever v is a star vertex. Thus, for example, the smoothing $\mathsf{S}_2 \to \mathsf{I}$ is allowed, but the inclusion $\mathsf{I} \to \mathsf{S}_n$ of a leg is not.

Proposition 4.9. *For any graph* Γ *and any* $U \subseteq V(\Gamma)$ *containing no isolated vertices, the inclusion* $\iota : S^U(\Gamma) \to S(\Gamma)$ *is a quasi-isomorphism.*

We omit a detailed proof, which can be seen, e.g., by a spectral sequence argument, filtering by polynomial degree in vertices and half-edges intersecting U.

Remark 4.10. The full Świątkowski complex has a canonical basis, while the reduced version lacks one in general. The reduction at (some subset of) the 1-valent vertices of a graph retains a canonical basis: if v is 1-valent, then $\widetilde{S}(v)$ is spanned by $\{\varnothing\}$. The corresponding reduced complex is the Świątkowski complex of a "graph" in which the 1-valent vertex has been deleted, leaving a half-open tail. All of our constructions and results can be made rigorous for such non-compact "graphs."

Our strategy in proving Theorem 4.5 will be to apply the tools developed earlier in the paper, especially Theorem 3.20. In order to do so, we must produce a decomposition, introduce Morse theory on the subdivisional configuration spaces of the pieces, and analyze the resulting Morse complexes.

CONSTRUCTION 4.11 (Canonical decomposition). Let Γ be a graph. Subdivide Γ by adding four vertices to each edge; denote the resulting graph by $\Gamma_\#$. There is a canonical graph morphism $\Gamma_\# \to \Gamma$ inducing a homeomorphism on configuration spaces.

Removing from each edge of Γ the open interval defined by the outer pair of added vertices produces a graph $\widetilde{\Gamma}_0 \cong \coprod_{v\in V} \mathsf{S}_{d(v)}$. On the other hand, for each edge, we have the closed interval defined by the inner pair of added vertices, and we let $\widetilde{\Gamma}_1 \cong E \times \mathsf{I}$ be the union of these closed intervals. Clearly, we have

$$\Gamma_\# \cong \widetilde{\Gamma}_0 \coprod_{A_1\times\{0\}} (A_1 \times \mathsf{I}) \coprod_{A_1\times\{1\}} \widetilde{\Gamma}_1,$$

where A_1 is a finite set with $|A_1| = 2|E|$.

We define the *canonical decomposition* of $\Gamma_\#$ (abusively, of Γ) by choosing the full poset of subdivisions $\mathrm{SD}(\Gamma_\#)$, which is filtered and convergent.

The main ingredient in the proof of Theorem 4.5 is the following.

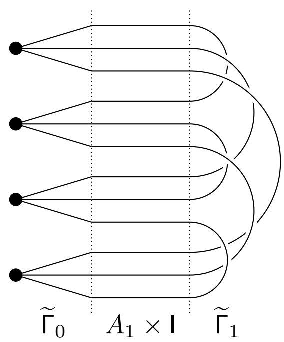

Figure 5: The canonical decomposition of the complete graph K_4

PROPOSITION 4.12. *Let Γ be a graph, and regard B^{SD} as a local invariant on the canonical decomposition of Γ. There is an isotopy invariant local flow on*

B^{SD} *such that*

$$\begin{aligned} I(A_1 \times \mathsf{I}) &\cong \mathbb{Z}[E] \otimes \mathbb{Z}[E] \\ I(\Gamma_0) &\cong \bigotimes_{v \in V} S^{\curlywedge}(\mathsf{S}_{d(v)}) \\ I(\Gamma_1) &\cong \mathbb{Z}[E] \end{aligned}$$

as associative algebras and modules, respectively.

The proof of this result amounts to constructing a couple abstract flows, calculating their Morse complexes, and checking various compatibilities. These tasks will occupy our attention in §4.3–4.4 below, but the idea behind the constructions is very simple. Intuitively, we define a flow on a star by pulling the cone point up, allowing points to flow down the legs, and we define a flow on an interval by allowing points to flow according to some fixed orientation. The reader who finds this heuristic sufficiently convincing may skip ahead to the computations.

For now, we deduce the following:

Proof of Theorem 4.5, construction of isomorphism. Applying Theorem 3.20 with the local flow of Proposition 4.12 produces an isomorphism

$$H_*(B(\Gamma)) \cong H_*\left(\bigotimes_{v\in V} S^{\curlywedge}(\mathsf{S}_{d(v)}) \overset{\mathbb{L}}{\bigotimes_{\mathbb{Z}[E]\otimes\mathbb{Z}[E]}} \mathbb{Z}[E] \right).$$

Since $\bigotimes_{v\in V} S^{\curlywedge}(\mathsf{S}_{d(v)})$ is a free $\mathbb{Z}[E]\otimes\mathbb{Z}[E]$-module, the derived tensor product is computed by the ordinary tensor product, and the proof is complete upon noting the canonical isomorphism

$$\bigotimes_{v\in V} S^{\curlywedge}(\mathsf{S}_{d(v)}) \bigotimes_{\mathbb{Z}[E]\otimes\mathbb{Z}[E]} \mathbb{Z}[E] \cong S(\Gamma).$$

□

4.3 Interval and star flows

In this section, we construct the abstract flows (in the sense of §3.3) that form the building blocks of the local flows alluded to in Proposition 4.12. These flows are inspired by the discrete flows of [FS05] and could be constructed in the same manner, but, in working locally, we are able to avoid the machinery of discrete Morse theory.

Definition 4.13 (Interval flow). Let $\mathsf{I} \to \mathsf{I}'$ be a subdivision. We fix an orientation of I' and define an order on the vertices by declaring the negative direction to be the direction of decrease. Define an endomorphism π of $C(B_k^{\square}(\mathsf{I}'))$ by declaring that π

1. takes any (positively signed) generator which is a set of k many 0-cells to the (positively signed) set of the k least 0-cells and

2. takes any generator which is a set containing a 1-cell to zero.

Using that the complex $B_k^{\square}(\mathsf{I}')$ is either contractible or empty, depending on whether k is greater than the number of 0-cells, the following result is easily verified.

LEMMA 4.14. *The interval flow is an abstract flow on* $C(B_k^{\square}(\mathsf{I}'))$.

DEFINITION 4.15 (Star flow). Let $\mathsf{S}_n \to \mathsf{S}'_n$ be a subdivision. For convenience, order the vertices in the ith leg by declaring the direction away from the star vertex v to be the direction of decrease; call them $v_{i,0}, \ldots, v_{i,N_i}$, with v_0 the least vertex and $v_{N_i} = v$ the star vertex. Let $e_{i,j}$ be the 1-cell containing $v_{i,j-1}$ and $v_{i,j}$. Given a tuple $(r_1, \ldots, r_n)$ with $0 \leq r_i \leq N_i$, write $V_{\vec{r}}$ for the set of vertices containing the r_i least vertices in the ith leg. Define an endomorphism π of $C(B_k^{\square}(\mathsf{S}'_n))$ by declaring that π

1. takes any (positively signed) set of 0-cells not including v with r_i 0-cells in the ith leg to the (positively signed) set $V_{\vec{r}}$,

2. takes any (positively signed) set of 0-cells including v and r_i additional 0-cells in the ith leg to the (positively signed) union of $V_{\vec{r}}$ and $\{v\}$,

3. takes any set containing a 1-cell not of the form e_{j,N_j} to zero, and

4. takes any set containing e_{j,N_j} and r_i additional (positively signed) 0-cells in the ith leg (necessarily not including v) to the sum

$$\sum_{r=r_j+1}^{N_j} V_{\vec{r}} \cup \{e_{j,r}\}.$$

See Figure 6.

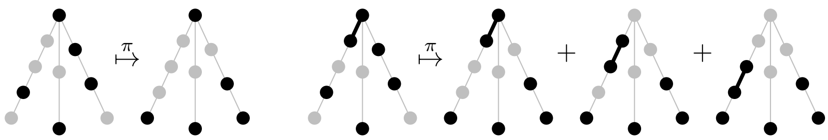

(a) 0-cell containing the star vertex (b) 1-cell containing a central edge

Figure 6: Examples of two of the cases for star flow (Definition 4.15)

LEMMA 4.16. *The star flow is an abstract flow on* $C(B_k^{\square}(\mathsf{S}'_n))$.

Proof. Idempotence is immediate except in the last case. There the image of the cell containing e_{j,N_j} and r_i vertices in the ith leg is the sum of the cell $V_{\vec{r}} \cup \{e_{j,N_j}\}$ with several cells in the kernel of π. The map π is a chain map by inspection (it suffices to check 2-cell and 1-cell generators).

Let us see that π is a quasi-isomorphism. For brevity, we write C for $C(B_k^{\square}(\mathsf{S}'_n))$ and I for its Morse complex. Consider the subcomplex A of C spanned by generators (i.e., sets of cells of $B_k^{\square}(\mathsf{S}'_n)$) which do not contain any 1-cell intersecting v. Since $\pi(A) \subseteq A$, there is a commuting diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & C & \longrightarrow & C/A & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & A\cap I & \longrightarrow & I & \longrightarrow & I/(A\cap I) & \longrightarrow & 0 \end{array}$$

of exact sequences. The quotient complex C/A is linearly isomorphic to the subspace $B \subseteq C$ spanned by generating cells containing an edge intersecting v, but, in the quotient, the differential ignores the special edge. Similarly, $I/(A \cap I) \cong (\pi(B), 0)$. It follows that $A \xrightarrow{\pi} (A \cap I)$ and $C/A \xrightarrow{\pi} I/(A \cap I)$ are both quasi-isomorphisms, since the domains are either contractible or empty in every weight, while the codomains have no differentials. The claim follows by the five lemma. □

Lemma 4.17. *There is a canonical chain map $I(B_k^{\square}(\mathsf{S}'_n)) \to S^{\curlywedge}(\mathsf{S}_n)_k$ that is an isomorphism if S'_n has at least $k+1$ vertices in each leg.*

Proof. A basis for the weight k subcomplex of $S^{\curlywedge}(\mathsf{S}_n)$ is given by the set of elements of the form $e_1^{r_1} \cdots e_n^{r_n} \otimes x$, where $x \in \{\varnothing, v, h_1, \ldots, h_n\}$, the r_j are non-negative integers, and $\sum r_j + \mathrm{wt}(x) = k$.
The coimage of π in the star flow is a quotient of the set of cells containing either the 0-cell v, the 1-cell e_{j,N_j}, or no cell containing v, along with r_j additional 0-cells in the jth leg for each j. The relation identifies two such configurations with the same central configuration and the same r_j for each j but different choices of 0-cells away from the center vertex. Then we send such a configuration to the element $e_1^{r_1} \cdots e_n^{r_n} \otimes x$ where x is:

1. the element v if c contains the star vertex v,
2. the element h_i if c contains the edge e_{j,N_j}, and
3. the element $\varnothing$ otherwise.

This map is prima facie injective, and it is surjective if S'_n has the hypothesized number of vertices in each leg. It is straightforward to verify that this map is also a chain map. □

4.4 Compatibilities

In this section, we check that the abstract flows constructed in the previous section are compatible with inclusion and subdivision, completing the proof of Proposition 4.12 and thereby of the isomorphism statement of Theorem 4.5.
Let $\Gamma \to \Gamma'$ be a subdivision of graphs, with Γ either the interval graph I or the star graph S_n, and let $\Xi' \subseteq \Gamma'$ be a subgraph that is a subdivision of a disjoint

union Ξ of intervals and stars; thus, $\Xi \to \Gamma$ is a graph morphism. We equip $B_k^{\square}(\Gamma')$ with either the interval flow or the star flow, depending on the case in consideration. Through the isomorphism of Lemma 2.11, the configuration complex $B_k^{\square}(\Xi')$ also inherits an abstract flow, which depends on choices of star points, since subdivided intervals, 1-stars, and 2-stars are intrinsically indistinguishable, and on choices of orientation for the interval components.

Lemma 4.18. *The inclusion* $i : B_k^{\square}(\Xi') \to B_k^{\square}(\Gamma')$ *is flow compatible provided that any component of* Ξ' *containing a star vertex of* Γ' *carries the star flow at that vertex. Moreover, in this case, the following diagram commutes:*

$$\begin{array}{ccc} I(B_k^{\square}(\Xi')) & \xrightarrow{I(B_k^{\square}(i))} & I(B_k^{\square}(\Gamma')) \\ \downarrow & & \downarrow \\ S^{\curlywedge}(\Xi)_k & \xrightarrow{S^{\curlywedge}(\Xi\to\Gamma)} & S^{\curlywedge}(\Gamma)_k. \end{array}$$

Proof. We prove only the case when Γ is a star graph and Γ' is equipped with the star flow (the interval case is easier), focusing on the verification of flow compatibility. Write π_Ξ for the abstract flow on $B_k^{\square}(\Xi')$ and π_Γ for the abstract flow on $B_k^{\square}(\Gamma')$. Let c be a cell of $B_k^{\square}(\Xi')$, written as a symmetric product of cells of Ξ'.

If c contains an edge that is not adjacent to the star vertex of Γ', then so does $i(c)$. Then by Definitions 4.13 and 4.15, $\pi_\Gamma(i(c)) = 0 = \pi_\Gamma(i(\pi_\Xi(c)))$. On the other hand, if c contains no such edge, then by the same definitions and inspection, $\pi_\Xi(c)$ is the sum of a cell c', which is obtained by moving the vertices of c into their minimal positions in the relevant components of Ξ', with a linear combination of cells, each containing an edge not adjacent to the star vertex of Γ'. The image of these latter cells are all in the kernel of π_Γ, so $\pi_\Gamma(i(\pi_\Xi(c)) = \pi_\Gamma(i(c'))$. But the same characterization also shows that $\pi_\Gamma(i(c')) = \pi_\Gamma(i(c))$, since both are obtained by moving the same number of vertices in each leg of Γ' into their minimal positions.

The commutativity claim is essentially immediate from what has already been said and the description of the isomorphism in Lemma 4.17. □

Similar considerations apply in the case of a subdivision:

Lemma 4.19. *For subdivisions* $\mathsf{I}' \to \mathsf{I}''$ *and* $\mathsf{S}'_n \to \mathsf{S}''_n$ *of the interval and the star graph* S_n *respectively, the induced subdivisional embeddings* $B_k^{\square}(\mathsf{I}') \to B_k^{\square}(\mathsf{I}'')$ *and* $B_k^{\square}(\mathsf{S}'_n) \to B_k^{\square}(\mathsf{S}''_n)$ *are flow compatible, and the following diagrams commute:*

$$\begin{array}{ccc} I(B_k^{\square}(\mathsf{I}')) & \longrightarrow & I(B_k^{\square}(\mathsf{I}'')) \\ \downarrow & & \downarrow \\ \mathbb{Z} & = & \mathbb{Z} \end{array} \qquad \begin{array}{ccc} I(B_k^{\square}(\mathsf{S}'_n)) & \longrightarrow & I(B_k^{\square}(\mathsf{S}''_n)) \\ \downarrow & & \downarrow \\ S^{\curlywedge}(\mathsf{S}_n)_k & = & S^{\curlywedge}(\mathsf{S}_n)_k. \end{array}$$

Proof of Proposition 4.12. By Lemma 4.19, the interval and star flows determine subdivisional flows on any disjoint union of intervals and stars equipped with any subdivisional structure. By the same result and Lemma 4.18, any embedding among such is flow compatible if the preimage of every star point is a star point. We use the star flow on component type basics in Γ_0 with the vertices of Γ as star vertices and the interval flow for all other basics. Since these subdivisional flows respect the isomorphisms $B^{\mathrm{SD}}(U \amalg V) \cong B^{\mathrm{SD}}(U) \times B^{\mathrm{SD}}(V)$, we obtain a local flow on the local invariant B^{SD}, isotopy invariant by Lemma 4.18. By inspection in the interval case and Lemma 4.17, there is an identification of the Morse complex $I(B^{\mathrm{SD}}(U)) \cong S^{\curlywedge}(U)$ which is natural for inclusions among disjoint unions of basics by Lemmas 4.18 and 4.19. Isotopy invariance and the identification of the algebra and module structures follow. □

4.5 Naturality for graph morphisms

In this section, we show that the isomorphism of Theorem 4.5 is functorial. To use the naturality clause of Theorem 3.20 we will show that every graph morphism lifts to a flow compatible map of local invariants. It suffices to do so for graph embeddings and smoothings separately.

If $f : \Gamma_1 \to \Gamma_2$ is a graph embedding, then every edge of Γ_1 is mapped homeomorphically to an edge of Γ_2, so we may choose auxiliary vertices so that there is a commuting diagram of graph morphisms:

$$\begin{array}{ccc} \Gamma_1 & \xrightarrow{f} & \Gamma_2 \\ \uparrow & & \uparrow \\ (\Gamma_1)_\# & \longrightarrow & (\Gamma_2)_\#. \end{array}$$

Then from the definition of the canonical decomposition, we have a commuting diagram of graph embeddings, and, in particular, of subdivisional spaces, which is to say a map of decompositions:

$$\begin{array}{ccccc} \coprod_{v\in V(\Gamma_1)} \mathsf{S}_{d(v)} & \longleftarrow & \coprod_{e\in E(\Gamma_1)} \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \coprod_{e\in E(\Gamma_1)} \mathsf{I} \\ \downarrow & & \downarrow & & \downarrow \\ \coprod_{v\in V(\Gamma_2)} \mathsf{S}_{d(v)} & \longleftarrow & \coprod_{e\in E(\Gamma_2)} \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \coprod_{e\in E(\Gamma_2)} \mathsf{I}. \end{array}$$

It follows from Lemmas 4.18 and 4.19 that the identity map on B^{SD} is flow compatible when regarded as a map between the induced local invariants on these two decomositions. Since every basic is a disjoint union of intervals and stars, the same lemmas show that the induced map at the level of Morse complexes coincides with the induced map on Świątkowski complexes. Thus, naturality holds for graph embeddings.

The case of a smoothing reduces to the following scenario. Fix an edge $e_0 \in \mathsf{\Gamma}$, and let $\mathsf{\Gamma}'$ be the graph obtained from $\mathsf{\Gamma}$ by adding a bivalent vertex v_0 to e_0, subdividing it into e_1 and e_2. There is a smoothing $f : \mathsf{\Gamma}' \to \mathsf{\Gamma}$, and every smoothing is a composite of isomorphisms and smoothings of this kind. This smoothing does not respect the canonical decompositions of $\mathsf{\Gamma}'$ and $\mathsf{\Gamma}$, but it does respect a different pair of decompositions $\mathcal{E}'$ and $\mathcal{E}$, depicted in the following diagram and Figure 7:

$$\begin{array}{ccccccccc}
\coprod_{v_0 \neq v \in V'} \mathsf{S}_{d(v)} & \longleftarrow & \coprod_{e \in E} \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \coprod_{e \in E'} \mathsf{I} & \longleftarrow & \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \mathsf{S}_2 \\
\| & & \| & & \| & & \| & & \downarrow \\
\coprod_{v \in V} \mathsf{S}_{d(v)} & \longleftarrow & \coprod_{e \in E} \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \coprod_{e \in E'} \mathsf{I} & \longleftarrow & \mathrm{pt} \amalg \mathrm{pt} & \longrightarrow & \mathsf{I}.
\end{array}$$

As before, Lemma 4.18 guarantees that the corresponding map of local invari-

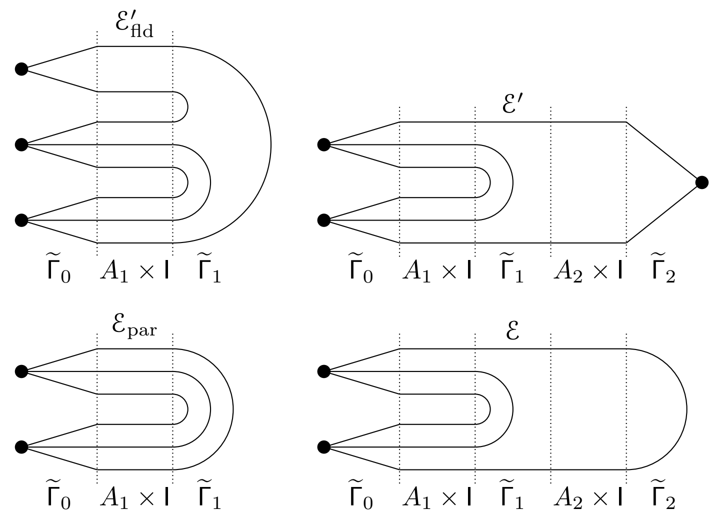

Figure 7: The canonical decomposition (labelled $\mathcal{E}'_{\mathrm{fld}}$) and the decomposition $\mathcal{E}'$ of the theta graph Θ_3 with an additional bivalent vertex are depicted in the first row. The canonical decomposition (labelled $\mathcal{E}_{\mathrm{par}}$) and the decomposition $\mathcal{E}$ of the theta graph Θ_3 itself are depicted in the second row. The notation used is that of Appendix B; compare Figure 11.

ants is flow compatible; moreover, Theorem 3.20 implies that the right-hand cell commutes in the diagram below. The left-hand cell commutes by inspection. Since the commuting of the outer square is what is to be proven, it will then suffice to verify that the two remaining triangles also commute.

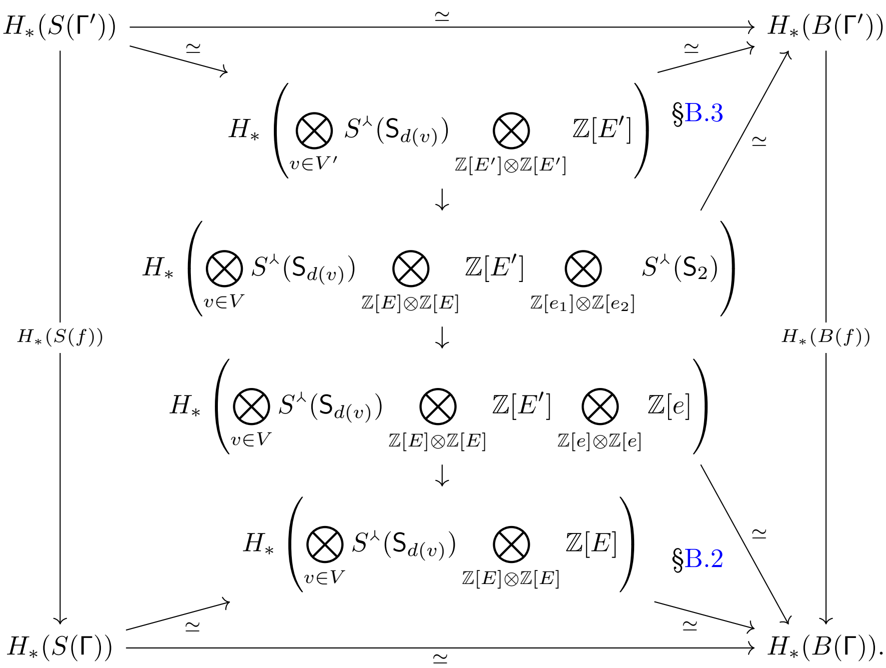

This question of comparing different decompositions of a fixed complex is the subject of Appendix B. The main tool there, Proposition B.5, is used twice. In §B.2, it is used to establish a compatibility which is applicable to $\mathcal{E}$ and the canonical decomposition of Γ. This first compatibility guarantees the commutativity of the lower triangle provided that every inclusion among disjoint unions of basics for the two decompositions is flow compatible. Then, in §B.3, the proposition is used to establish a compatibility applicable to $\mathcal{E}'$ and the canonical decomposition of Γ'. This second compatibility guarantees the commutativity of the upper triangle provided the corresponding inclusions are flow compatible. As before, these compatibilities are a direct consequence of Lemma 4.18.

5 Homology of graph braid groups

We begin our computational study of the homology of configuration spaces of graphs in earnest. Highlights include Corollary 5.16, an exact sequence associated to the removal of a vertex; Proposition 5.21, which asserts that multiplication by any fixed edge is injective; Proposition 5.25, which identifies the homology in top degree in the case of a trivalent graph; Proposition 5.30, a full computation in the case of the complete graph on four vertices; and Appendix C, which gives a streamlined derivation of the characterization of the first homology due to [KP12]. Except for this last, these results appear to be new.

5.1 Graphs, star and loop classes, and relations

Since a graph embedding induces a map at the level of Świątkowski complexes, some homology classes originate in subgraphs. Moreover, the homology of the configuration spaces of a graph is constrained by relations originating in its subgraphs. In this section, we acquaint ourselves with a few useful generators and relations of this kind. The generators we discuss are present in [HKRS14, Section 4], as is the Q-relation of Definition 5.9 (present as their equation (9) and Lemma 3).

Recall that the star graph S_n is the cone on the set $\{1,\ldots,n\}$. This graph has $n+1$ vertices: a central vertex v_0 of valence n and n vertices $\{v_1,\ldots,v_n\}$ of valence 1. Its half-edges E are $\{e_1,\ldots,e_n\}$, and it has $2n$ corresponding half-edges H, of which n of these, $h_1,\ldots,h_n$, are at v_0 and one, h_i', is at v_i for $i\neq 0$.

We will also consider several other graphs in this section.

Definition 5.1. The *cycle graph* C_n is a topological circle equipped with n bivalent vertices and n edges. The *lollipop graph* L_n is obtained by attaching an extra edge e_0 to the cycle graph C_n at one half-edge. The *theta graph* Θ_n is the topological suspension (double cone) on n points, with vertices the cone points v_1 and v_2.

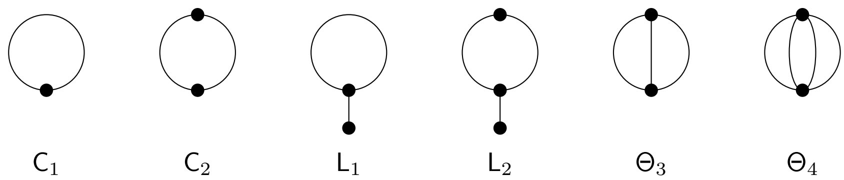

Figure 8: Cycle graphs, lollipop graphs, and theta graphs

The reduced Świątkowski complex of the star graph $\widetilde{S}(\mathsf{S}_n)$ is concentrated in degrees 0 and 1 with a single differential

$$\mathbb{Z}[E]\langle h_{12},\ldots,h_{1n}\rangle \xrightarrow{\partial} \mathbb{Z}[E]\langle\varnothing\rangle.$$

Lemma 5.2. *1. The homology group $H_1(B(\mathsf{S}_3))$ is freely generated as a $\mathbb{Z}[E]$-module by a single class α in $H_1(B_2(\mathsf{S}_3))$.*

2. The homology group $H_1(B(\mathsf{C}_n))$ is generated as a $\mathbb{Z}[E]$-module by a single class γ in $H_1(B_1(\mathsf{C}_n))$ subject to the relations $e_i\gamma = e_j\gamma$.

Proof. For S_3, it is a straightforward to see that $H_0(B(\mathsf{S}_3))$ is rank one in each weight. By observation, the chain $a \coloneqq a_{123} = e_1h_{23} + e_2h_{31} + e_3h_{12}$ is closed.

Checking the Euler characteristic shows that $\mathbb{Z}[E]\langle a\rangle$ has the correct rank and thus is the entirety of the kernel of ∂.

For C_1, the complex $\widetilde{S}(\mathsf{C}_1)$ has no differential and its degree one subspace is isomorphic to the module described. There is a (non-unique) smoothing from C_n to C_1; the induced map on Świątkowski complexes identifies every edge of C_n with the unique edge of C_1. □

The sign of the generating class α in $H_1(B_2(\mathsf{S}_3))$ depends on the choice of ordering of the half-edges. We will use the notation $\alpha_{123} = \alpha_{231} = \alpha_{312}$ for α with the convention employed here and $\alpha_{132} = \alpha_{321} = \alpha_{213}$ for the generator with the opposite sign.

Definition 5.3. Let $\mathsf{S}_3 \to \Gamma$ be a graph morphism. We call the image of the class α in $H_1(B_2(\Gamma))$ a *star class* and a representing cycle of the form above a *star cycle*.

Let $\mathsf{C}_n \to \Gamma$ be a graph morphism. We call the image of the class γ in $H_1(B(\Gamma))$ a *loop class* and a representing cycle a *loop cycle*.

A star class depends only on the isotopy class of the graph morphism which induces it. Note also that we have such a morphism whenever there are three distinct half-edges at a vertex; it is not necessary that the corresponding edges be distinct.

Lemma 5.4. *Star and loop classes are non-trivial.*

Proof. There is a natural homomorphism $\sigma_\Gamma : H_1(B_2(\Gamma)) \to \mathbb{Z}/2\mathbb{Z}$, which is induced on Abelianizations by the homomorphism from the braid group $\pi_1(B_2(\Gamma))$ recording the permutation of the endpoints of a braid. By naturality, evaluating σ_Γ on a star class gives the same answer as evaluating σ_{S_3} on a generator of $H_1(B_2(\mathsf{S}_3))$. Since any cycle representing such a generator interchanges the two points of the configuration, σ_Γ takes the value 1 on this generator.

Loop classes live in $H_1(B_1(\Gamma))$ which is naturally isomorphic to $H_1(\Gamma)$; a loop in Γ is never trivial. □

Star classes play a pivotal role in the remainder of the paper. For example, we have the following result.

Lemma 5.5. *The $\mathbb{Z}[E]$-module $H_1(B_k(\mathsf{S}_n))$ is generated by star classes.*

Similar statements go back at least to [FS05, FS12].

Proof. We proceed by induction on n. Since S_n is topologically an interval for $n \in \{1,2\}$, the claim holds trivially in these cases, and the case $n = 3$ follows from Lemma 5.2, so we may assume that $n \geq 4$.

Any degree 1 element $a \in S^{\partial}(\mathsf{S}_n)$ can be written in the form $a = \sum_{i=1}^n p_i h_i$ with $p_i \in \mathbb{Z}[E]$, and imposing the condition $\partial a = 0$ yields the two equations

$$\partial a = \sum_{i=1}^{n} p_i(e_i - v) = 0 \iff \begin{cases} \sum_{i=1}^n p_i = 0; \\ \sum_{i=1}^n e_i p_i = 0. \end{cases} \tag{1}$$

For each $i \leq n-2$, we write $p_i = (e_n - e_{n-1})p'_i + r_i$, where r_i does not involve the variable e_n. Now we can rewrite a partially in terms of the star cycles $a_{i,n-1,n}$ as follows:

$$a = \sum_{i=1}^{n-2}(p'_i a_{i,n-1,n} + r_i h_i) + q_{n-1}h_{n-1} + q_n h_n,$$

where

$$q_{n-1} \coloneqq p_{n-1} - \sum_{i=1}^{n-2}(e_i - e_n)p'_i, \qquad q_n \coloneqq p_n - \sum_{i=1}^{n-2}(e_{n-1} - e_i)p'_i.$$

Write q_{n-1} and q_n as

$$q_{n-1} = e_n q'_{n-1} + r_{n-1}, \qquad q_n = e_n q'_n + r_n,$$

where r_n and r_{n-1} do not involve the variable e_n. Then considering terms involving e_n in the equations (1) using the fact that $\partial a_{i,n-1,n} = 0$, we have

$$\begin{cases} e_n(q'_{n-1} + q'_n) = 0; \\ e_{n-1}q'_{n-1} + e_n q'_n + r_n = 0 \end{cases} \iff (e_n - e_{n-1})q'_n + r_n = 0.$$

Since r_n does not involve the variable e_n, it follows that $q'_n = r_n = 0$, whence $q_n = 0$. We further conclude that $q'_{n-1} = 0$, and so q_{n-1} does not involve the variable e_n; therefore, $a - \sum_{i=1}^{n-2} a_{i,n-1,n}$ does not involve e_n or h_n and so must lie in the image of the map $S^{\partial}(-)$ induced by the inclusion $\mathsf{S}_{n-1} \to \mathsf{S}_n$ that misses the nth leg. The inductive hypothesis now completes the proof. □

PROPOSITION 5.6. *If Γ is connected, then $H_1(B(\Gamma))$ is generated as a $\mathbb{Z}[E]$-module by star classes and loop classes.*

Remark 5.7. Again this is more or less implicit in the work of Farley and Sabalka [FS05, FS12]. We provide a proof using our language and methods.

Proof. Assume first that Γ is a tree. In this case the claim follows from Proposition 5.22, Lemma 5.5, and induction on the number of vertices of Γ of valence at least three. In the general case, subdivide each edge of Γ into three edges, calling the resulting graph Γ'. Each edge of Γ corresponds to three edges of Γ'; let E_{mid} be the set of edges of Γ' none of whose vertices are vertices of Γ. Let E_0 be a subset of E_{mid} whose complement is a spanning tree T of Γ'. For $e \in E_0$, write $\Gamma^{[e]}$ for the disjoint union of the unique cycle subgraph C_e contained in $e \cup \mathsf{T}$ with a disjoint edge for each edge of $E_{\mathrm{mid}} \backslash E(\mathsf{C}_e)$. There are canonical graph embeddings of T and C_e into Γ' which induce graph morphisms from T and C_e into Γ. Write $\widehat{S}(\Gamma) \coloneqq \widetilde{S}(\mathsf{T}) \oplus \bigoplus_{e \in E_0} \widetilde{S}(\Gamma^{[e]})$; then we have a map of differential graded $\mathbb{Z}[E]$-modules

$$\phi : \widehat{S}(\Gamma) \to \widetilde{S}(\Gamma)$$

induced by these graph morphisms. It is clear that $H_1(B(\mathsf{\Gamma}^{[e]}))$ is generated as a $\mathbb{Z}[E]$-module by its unique loop class for each $e \in E_0$, and we have already shown that $H_1(B(\mathsf{T}))$ is generated by star classes; therefore, it will suffice to show that ϕ induces a surjection on H_1.
For $e \in E_0$, write e' and e'' for the unique pair of edges of T with $\phi(e) = \phi(e') = \phi(e'')$. At the chain level, $\phi|_{\widetilde{S}_1(\mathsf{T})}$ is surjective with kernel the $\mathbb{Z}[E]$-span of the set $\{(e' - e'') : e \in E_0\}$. Let $c_e \in \widetilde{S}(\mathsf{C}_e)$ be a cycle representing a loop class and choose a degree 1 element $b_e \in \widetilde{S}(\mathsf{T})$ with $\phi(b_e) = \phi(c_e)$ and $\partial b_e = \pm(e' - e'')$. Then $\phi(b_e - c_e) = 0$ and $\partial(b_e - c_e) = \pm(e' - e'') \in \widetilde{S}(\mathsf{T})$.
Now, suppose we are given $b \in \widetilde{S}(\mathsf{T})$ with $\phi(\mathsf{T})$ closed. Then

$$\partial b = \sum_{e \in E_0} p_e(e' - e''),$$

and we set

$$b' := b + \sum_{e \in E_0} \mp p_e(b_e - c_e) \in \widehat{S}(\mathsf{\Gamma}).$$

Then $\phi(b') = \phi(b)$ and b' is closed, as desired. □

Now we turn to relations.

Lemma 5.8. *1. The homology group $H_0(B(\mathsf{S}_2))$ is generated by the class of the empty configuration $\varnothing$ subject to the relation $e_1\varnothing = e_2\varnothing$.*

2. In $H_1(B(\mathsf{S}_4))$, the star classes satisfy the relation

$$e_1\alpha_{234} - e_2\alpha_{341} + e_3\alpha_{412} - e_4\alpha_{123} = 0.$$

3. In $H_1(B(\mathsf{L}_n))$, the loop and star classes satisfy the relation

$$(e - e_0)\gamma = \alpha$$

where e is an edge of the cycle subgraph.

4. Let the edges of Θ_3 be numbered from 1 to 3 and likewise for the half-edges at v_1 and at v_2. Then the star class α_{123} at v_1 and the star class α_{321} at v_2 are equal.

Proof. 1. The cokernel of the differential is generated freely by $(e_1 - e_2)$.

2. The claim follows by expansion of the star cycles a_{ijk}.

3. It suffices to verify the claim for L_1 where it is already true for the (unique) chain level representatives using the reduced Świątkowski complex.

4. The chain $h_{12} \otimes h'_{13} - h_{13} \otimes h'_{12}$ in $S_2(\Theta_3)$ bounds the difference between the corresponding star cycles.

□

DEFINITION 5.9. Let Γ be a graph.

1. Let S_2 to Γ be a graph morphism. We call the induced relation in $H_0(B(\Gamma))$ an *I-relation.*

2. Let $\mathsf{S}_4 \to \Gamma$ be a graph morphism. We call the induced relation on star classes in $H_1(B(\Gamma))$ an *X-relation.*

3. Let $\mathsf{L}_n \to \Gamma$ be a graph morphism. We call the induced relation on loop and star classes in $H_1(B(\Gamma))$ a *Q-relation.*

4. Let $\Theta' \to \Theta_3$ be a smoothing and $\iota : \Theta' \to \Gamma$ a graph embedding. Denoting the preimage of v_i in Θ' by v_i', we call the relation induced on the star classes at $\iota(v_1')$ and $\iota(v_2')$ in $H_1(B(\Gamma))$ a *Θ-relation.*

5. Let $\mathsf{C}_n \to \Gamma$ be a graph morphism. We call the induced relation (from Lemma 5.2) on loop classes in $H_1(B(\Gamma))$ an *O-relation.*

Repeated use of the I-relation implies the well-known fact that $H_0(B_k(\Gamma))$ is one-dimensional for Γ connected.
These atomic classes combine naturally.

DEFINITION 5.10. Consider a graph Γ_0 written as the disjoint union of n_1 cycle graphs and n_2 copies of S_3. Since the configuration spaces of the components all have torsion-free homology, there is a class β in $H_{n_1+n_2}(B_{n_1+2n_2}(\Gamma_0))$ corresponding to the tensor product of the loop classes in each cycle graph and the star classes in each star graph. If $\Gamma_0 \to \Gamma$ is a graph morphism, we call the image of β in $H_{n_1+n_2}(B_{n_1+2n_2}(\Gamma))$ the *external product* of the corresponding loop and star classes. We will use juxtaposition to indicate the external product—writing $\alpha_1\alpha_2$ for the external product of classes α_1 and α_2—and we caution the reader that this construction is only partially defined and so does not define a product on homology.

5.2 EULER CHARACTERISTIC

One calculation that requires no further tools is that of the Euler characteristic. This result has been known at least since [Gal01].

COROLLARY 5.11. *The Euler characteristic of $B_k(\Gamma)$ is given by*

$$\chi(B_k(\Gamma)) = \sum_{U \subseteq V} (-1)^{|U|} \binom{k - |U| + |E| - 1}{|E| - 1} \prod_{v \in U} (d(v) - 1).$$

Defining the Euler–Poincaré series of Γ to be the formal power series

$$P_\chi(\Gamma)(t) = \sum_{k=0}^{\infty} \chi(B_k(\Gamma)) t^k,$$

we have the following reformulation.

Corollary 5.12. *Let* Γ *be a graph. Then*

$$P_\chi(\Gamma)(t) = \prod_v \frac{1 + t(1 - d(v))}{(1-t)^{\frac{d(v)}{2}}}.$$

Proof. For $n \geq 0$, the formula holds for the graph G_{2n} that is the wedge of n circles. It also holds, for $m \leq n$ non-negative, for the graph $\mathsf{G}_{2m+1,2n+1}$, which is given by connecting the vertices of G_{2m} and G_{2n} by a single edge. Since the formula of Corollary 5.11 does not depend on which pairs of vertices share an edge, the claim follows. □

5.3 Spectral sequences and exact sequences

One particularly nice class of tools afforded by the Świątkowski complex consists of spectral sequences that arise by decomposing the differential into a bicomplex. In general these facilitate the reduction of computations of $H_*(B(\Gamma))$ to computations for simpler graphs. Choosing our input data judiciously, we get an exact sequence which arises from deleting a vertex of the graph.

Lemma 5.13. *Let* J *be a set of half-edges of a graph* Γ*, and let* U *be the subset of vertices of* Γ *that half-edges of* J *are incident on. Then:*

1. *the differential* ∂ *of the Świątkowski complex decomposes into the sum of two commuting* $\mathbb{Z}[E]$*-linear differentials* $\partial_J + (\partial - \partial_J)$*, where* ∂_J *changes the number of generators containing half-edges in* J*, and*

2. *the (first) spectral sequence associated to this bicomplex collapses at the* $(|U|+1)$*st page, necessarily to the homology of the Świątkowski complex.*

Proof. For the first statement, it suffices to note, first, that ∂ can be rewritten as the sum $\partial = \sum \partial_h$ over all half-edges, and, second, that the operators ∂_h individually square to zero and pairwise commute.
For the second statement, since the differential ∂_J lowers the number of half-edges in J in a homogeneous monomial by 1 and since there are necessarily between 0 and $|U|$ such half-edges in any homogeneous monomial, all differentials starting from page $(|U|+1)$ vanish identically. □

When J consists of all half-edges incident on a set of vertices, there is a particularly straightforward description of the E^1 page of this spectral sequence, or rather of the corresponding spectral sequence for the reduced Świątkowski complex.

Definition 5.14. Let v be a vertex of the graph Γ. Then we write Γ_v for the *vertex explosion* of Γ at v, that is, the graph obtained by

1. replacing the vertex v with $\{v\} \times H(v)$ and

2. modifying the attaching map for half-edges attached at v by letting such a half-edge h be attached at $v \times h$.

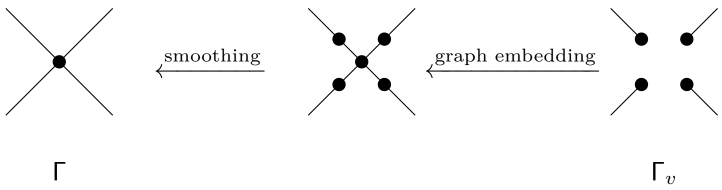

Figure 9: A local picture of vertex explosion along with an intermediate graph which admits a graph morphism from Γ_v and a smoothing to Γ.

There is a graph morphism from Γ_v to Γ which takes each edge to itself, takes the vertex $v \times h$ to $e(h)$, and takes each other vertex to itself. Defining this morphism requires choices (of precisely where in $e(h)$ to send $v \times h$) but the isotopy class of this graph morphism is unique. See Figure 9. More generally, if U is a set of vertices, then the isomorphism class of vertex explosion Γ_U at all vertices of U (by sequentially exploding the vertices of U) is well-defined and independent of the choice of order of explosion.

Lemma 5.15. *Suppose J is the set of all half-edges incident on a set U of vertices of a graph Γ. Then the spectral sequence of Lemma 5.13 is naturally defined for the reduced Świątkowski complex. In weight k, the entry $E^1_{p,q}$ in this spectral sequence for the reduced complex is (unnaturally) isomorphic to*

$$E^1_{p,q} \cong \bigoplus_{v_1,\ldots,v_q \in U} \bigoplus^{\prod(d(v_i)-1)} H_p(B_{k-q}(\Gamma_U)).$$

with differential as described in the proof.

Proof. The condition on J implies that we can understand the induced filtration on the reduced subcomplex by counting generators that are differences of half-edges incident on vertices in U, so that the bicomplex structure is natural in the reduced setting.
Moreover, the entry $E^0_{p,q}$ is isomorphic to the sum of copies of the reduced Świątkowski complex of Γ_U indexed by the choice of q half-edge differences at distinct vertices in U. The differential from $E^0_{p,q}$ to $E^0_{p-1,q}$ is the sum of indexed copies of the reduced Świątkowski differential of Γ_U, yielding the indicated homology groups.
By fixing a half-edge h^v_0 of each $v \in U$, we can index this double sum by sequences of q half-edges incident on distinct vertices of U, none of them h^v_0. Then the differential $d^1 : E^1_{p,q} \to E^1_{p,q-1}$ takes the signed sum over deleting a half-edge h from the sequence and multiplying the homology class by $e(h) - e(h^{v(h)}_0)$:

$$d^1(\alpha_{(h_1,\ldots,h_q)}) = \sum_{i=1}^{q} (-1)^{i-1} \left(e(h_i) - e(h_0^{v(h_i)}) \right) \alpha_{(h_1,\ldots,\widehat{h_i},\ldots,h_q)}.$$

□

In the computations of the rest of this paper, we will only use the simplest case of this spectral sequence, which degenerates at page E^2.

Corollary 5.16. *Fix a half-edge $h_0 \in H(v)$. There is a long exact sequence of differential bigraded $\mathbb{Z}[E]$-modules*

$$\cdots \to H_{n+1}(B_k(\Gamma_v)) \to H_{n+1}(B_k(\Gamma)) \longrightarrow \bigoplus_{h \in H(v)\setminus\{h_0\}} H_n(B_{k-1}(\Gamma)) \longrightarrow$$

$$\longrightarrow H_n(B_k(\Gamma_v)) \longrightarrow H_n(B_k(\Gamma)) \longrightarrow \bigoplus_{h \in H(v)\setminus\{h_0\}} H_{n-1}(B_{k-1}(\Gamma)) \to \cdots$$

The connecting homomorphism δ from $\bigoplus H_*(B(\Gamma_v)) \to H_*(B(\Gamma_v))\{-1\}$ is explicitly given by the formula

$$\delta\beta_h = (e(h) - e(h_0))\beta_h.$$

Proof. The spectral sequence of Lemma 5.15 in the case $U = \{v(h_0)\}$ collapses at E^2, yielding a long exact sequence with connecting homomorphisms the E^1 differentials, which are described in the proof of the lemma. □

We now use this spectral sequence to recover a homologico-algebraic version of our decomposition theorem (Theorem 3.20) for the following special type of decomposition.

Definition 5.17. Let Γ be a graph and k a non-negative integer. A *k-bridge decomposition* $\Gamma = \Gamma_1 \sqcup_L \Gamma_2$ consists of a set L of k distinct edges of Γ, called the *k-bridge*, and two subgraphs $\widetilde{\Gamma}_1$ and $\widetilde{\Gamma}_2$ of Γ such that

$$\Gamma = \widetilde{\Gamma}_1 \cup_{V_1} L \cup_{V_2} \widetilde{\Gamma}_2,$$

where $V_1 \cup V_2$ is a set of $2k$ distinct vertices. In this case, we write $\Gamma_j = \widetilde{\Gamma}_j \cup L$.

Proposition 5.18. *Let $\Gamma = \Gamma_1 \sqcup_L \Gamma_2$ be a graph with a k-bridge decomposition and $\mathbb{K}$ a field. There is a spectral sequence converging to $H(B(\Gamma), \mathbb{K})$ with E^2 page given by $\mathrm{Tor}^{\mathbb{K}[L]}(H(B(\Gamma_1), \mathbb{K}), H(B(\Gamma_2), \mathbb{K}))$. For k at least 2, this spectral sequence collapses at page k (for $k = 1$ it collapses at page 2).*

Proof. Subdivide each edge of the bridge and let J be the set of k new vertices. The spectral sequence of Lemma 5.15 for this vertex set J has a special property. Because each vertex is valence 2, the E^0 page has $E^1_{p,q} = \wedge^q H_p(B_{k-q}(\Gamma_1 \sqcup \Gamma_2))$. After passing to $\mathbb{K}$ coefficients and invoking the Künneth theorem, examination of the d^1 differential reveals that it is the Koszul differential which calculates the Tor groups of the statement of the proposition. Moreover, d^1 from $E^1_{p,k}$ to $E^1_{p,k-1}$ is injective by Proposition 5.21, which means that E^2 is concentrated in rows $q = 0$ to $q = k - 1$. □

Corollary 5.19. *Let $\Gamma = \Gamma_1 \sqcup_L \Gamma_2$ be a graph equipped with a 2-bridge decomposition. Then the associated graded of $H(B(\Gamma))$ is isomorphic to* $\mathrm{Tor}^{\mathbb{K}[L]}(H(B(\Gamma_1)), H(B(\Gamma_2)))$.

We defer consideration of the simpler situation of a 1-bridge decomposition to Proposition 5.22, which is an integral statement.

Remark 5.20. Under the quasi-isomorphism supplied by the decomposition theorem, the spectral sequence of Proposition 5.18 is identified with the Künneth spectral sequence.

5.4 Edge injectivity

We will make repeated use of the following result as a technical tool, but it is interesting in itself, and we think of it as a first step in an investigation of the $\mathbb{Z}[E]$-module structure enjoyed by the homology of the configuration spaces of Γ, in the tradition of the study of homological stability. This study is continued in detail in the sequel [ADCK]—see also [Ram18].

Proposition 5.21. *For any $e \in E$, multiplication by e is injective on $H_*(B(\Gamma))$.*

Proof. Any monomial $p(E, V)$ in the edges and vertices of Γ can be rewritten as a polynomial in variables e, $e - e'$, and $e - v$. Write U for the collection of generators other than e, which is in bijection with $E\backslash\{e\} \amalg V$. Then an arbitrary chain in $S(\Gamma)$ can be written in the form

$$b = \sum_{i=0}^{M} e^i b_i(U, H)$$

for some (graded commutative) polynomials b_i (not all polynomials are possible).

Let c be a cycle in $S(\Gamma)_k$, and suppose that $ec = \partial b$. The proof will be complete once we are assured that c is itself a boundary. To see this, we note that the differential, acting on $b_i(U, H)$, cannot introduce a factor of e in this basis; that is,

$$\partial b = \sum_{i=0}^{M} e^i b_i'(U, H),$$

where $b_i' = \partial b_i$. Then we have

$$\sum_{i=1}^{M} e^i c_{i-1}(U, H) = \sum_{i=0}^{M} e^i b_i'(U, H).$$

Thus, $\partial b_0 = b_0' = 0$, so $ec = \partial b = \partial(b - b_0)$. Since $b - b_0$ is divisible by e, we have $c = \partial(\frac{b - b_0}{e})$, as desired. ∎

With Proposition 5.21 in hand, we can describe more precisely the effect on $H_*(B(\Gamma))$ of cutting Γ into two disconnected components at a vertex.

Proposition 5.22. *Let Γ be a connected graph and v a bivalent vertex whose removal disconnects Γ. Write e_1 and e_2 for the edges at v. There is an isomorphism*

$$H_*(B(\Gamma)) \cong H_*(B(\Gamma_v))/e_1 \sim e_2$$

of $\mathbb{Z}[E]$-modules. In particular, for any field $\mathbb{K}$, we have

$$H_*(B(\Gamma), \mathbb{K}) \cong H_*(B(\Gamma'_v), \mathbb{K}) \bigotimes_{\mathbb{K}[e]} H_*(B(\Gamma''_v), \mathbb{K}),$$

where Γ'_v and Γ''_v are the connected components of Γ_v.

Proof. Applying Corollary 5.16 at the vertex v, we obtain the exact sequence

$$0 \to \widetilde{S}(\Gamma_v) \xrightarrow{\phi} \widetilde{S}(\Gamma) \xrightarrow{\psi} \widetilde{S}(\Gamma_v)[1]\{1\} \to 0.$$

We claim that the connecting homomorphism in the corresponding long exact sequence is injective. Indeed, suppose that $\delta\beta = 0$ for some $\beta \in H_i(B_k(\Gamma_v))[1]\{1\}$; then, applying the formula of Corollary 5.16, we have $(e_1 - e_2)\beta = 0$. Since Γ_v is disconnected, the homology of $B_k(\Gamma_v)$ is naturally graded by the number of points lying in Γ'_v; therefore, writing $\beta = \sum_{j=0}^k \beta_j$, we obtain the system of equations

$$\begin{aligned} e_1\beta_k &= e_2\beta_0 = 0 \\ e_1\beta_j &= e_2\beta_{j+1} \qquad (0 \le j < k). \end{aligned}$$

Applying Proposition 5.21 repeatedly now implies that $\beta_j = 0$ for every $0 \le j \le k$.

The statement over $\mathbb{Z}$ now follows by exactness and the formula for δ, and the statement over $\mathbb{K}$ follows. ☐

Remark 5.23. This situation has been considered before [MS17, Sections III and VI]. There, assuming their Conjecture VI.1, they prove the special case of Proposition 5.22 for homological degree two. To be more precise, their Theorem VI.2 provides a count of the rank of the homology group $H_2(B(\Gamma))$ in terms of $H_*(B(\Gamma'_v))$ and $H_*(B(\Gamma''_v))$. Proposition 5.21 implies that $H_*(B(\Gamma'_v))$ and $H_*(B(\Gamma''_v))$ are free $\mathbb{K}[e]$-modules, and their formula for the rank comes directly from the formula for the weighted rank of the tensor product of free $\mathbb{K}[e]$-modules (a side remark: there seems to be a typo in the range of the summation of their equation (24), which omits $\beta_1^{(0)}(\Gamma_1)\beta_1^{(n)}(\Gamma_2)$). It seems likely that the injectivity of our connecting homomorphisms used in the proof of Proposition 5.22 implies their Conjecture VI.1 (justifying the rest of their section VI), but their iterated Mayer–Vietoris decompositions differ enough from the "one-stage" decomposition of our vertex long exact sequence that it might take some work to translate between the two. Since Proposition 5.22 supercedes their Theorem VI.2, we have chosen not to further pursue such verification.

5.5 High degree homology for unitrivalent graphs

Call a graph *unitrivalent* if every vertex has valence either 3 or 1. In this section we apply the exact sequence of Corollary 5.16 to study $H_i(B(\Gamma))$ when Γ is unitrivalent and i is close to the number of trivalent vertices of Γ. According to our conventions, a graph may have self-loops (at most one self-loop per vertex by unitrivalency), and multiple edges connecting a given pair of vertices. In this subsection, unless otherwise specified, Γ is a unitrivalent graph and N denotes the number of trivalent vertices of Γ.

Construction 5.24. Suppose Γ is unitrivalent and has r self-loops. To each trivalent vertex v we associate a homology class $\beta_v \in H_1(B(\Gamma))$, well-defined up to sign: if v has a self-loop, then β_v is the corresponding loop class; otherwise, β_v is the star class at v. Taking the external product over the trivalent vertices of the β_v, we obtain a homology class $\beta \in H_N(B_{2N-r}(\Gamma))$, well-defined up to sign, called the *canonical class* of Γ.

Proposition 5.25. *For Γ unitrivalent, we have an isomorphism*

$$H_N(B(\Gamma)) \cong \mathbb{Z}[E]\beta.$$

Proof. Without loss of generality, we may assume Γ is connected. We proceed by induction on N with two base cases. The first base case is the 3-star, which was dealt with in §5.1. The second base case is the lollipop L_1. The reduced Świątkowski complex $\widetilde{S}(\mathsf{L}_1)$ is concentrated in degrees 0 and 1, with the explicit form

$$\widetilde{S}(\mathsf{L}_1) \cong \mathbb{Z}[e_0, e]\langle \varnothing, h_{01}, h_{02}\rangle$$

with differential $\partial(h_{01}) = \partial(h_{02}) = e_0 - e$. By inspection, the kernel of the differential in degree 1 is spanned by $\mathbb{Z}[e_0, e](h_{01} - h_{02})$.
Now assume that $N \geq 2$, fix a trivalent vertex v, and assume that the claim is known for Γ_v. Applying Corollary 5.16 at v, we obtain the exact sequence

$$\cdots \to H_N(\widetilde{S}(\Gamma_v)) \to H_N(\widetilde{S}(\Gamma)) \to \bigoplus_{h\in H(v)\setminus\{h_0\}} H_{N-1}(\widetilde{S}(\Gamma_v))\{1\} \xrightarrow{\delta} \cdots$$

Since Γ_v has $N-1$ trivalent vertices, the first term vanishes, so the homology group of interest is the kernel of δ. The inductive hypothesis identifies the domain of δ as

$$\left(\bigoplus_{h\in H(v)\setminus\{h_0\}} \mathbb{Z}[E(v)]\right) \otimes \mathbb{Z}[E\setminus E(v)]\beta',$$

where β' is the canonical class of Γ_v and δ acts on the first factor. Denoting by $\Gamma^{(v)}$ the subgraph containing v and all of its edges (either a star or a lollipop), this first factor is naturally identified with the degree 1 component of $\widetilde{S}(\Gamma^{(v)})$. Under this identification, it is immediate from the formula of Corollary 5.16 that the condition of lying in the kernel of δ is precisely the condition of lying in the kernel of ∂ from degree one to degree zero of $\widetilde{S}(\Gamma^{(v)})$. Since $\Gamma^{(v)}$ has only one vertex,

$$\ker \partial \cong H_1(B(\Gamma^{(v)})) \cong \mathbb{Z}[E(v)]\beta_v,$$

and the claim follows. ☐

Remark 5.26. For a general graph (not necessarily unitrivalent), it is straightforward to show that $H_N(B_k(\Gamma)) = 0$ for $k < N$. Similar methods to those used in Proposition 5.25 show, for instance, that if Γ has no bivalent vertices or self-loops and at least one vertex, that $H_N(B_N(\Gamma)) = 0$ as well. Any further sharpening of this result must take into account, e.g., that $H_2(B_3(\Theta_4))$ is non-zero.

Similar inductive arguments using Corollary 5.16 may be used to demonstrate the following two results. Recall that a graph is simple if it has no self-loops and no multiple edges.

Proposition 5.27. *Suppose Γ is simple and unitrivalent. There is an isomorphism*

$$H_{N-1}(B(\Gamma)) \cong \bigoplus_{d(v)=3} \mathbb{Z}[E]\hat{\beta}_v/(e\hat{\beta}_v - e'\hat{\beta}_v : e, e' \in E(v))$$

of $\mathbb{Z}[E]$-modules, where $\hat{\beta}_v \in H_{N-1}(B_{2N-2}(\Gamma))$ is the external product of β_w for $w \neq v$.

Proposition 5.28. *If Γ is simple and unitrivalent, then $H_{N-2}(B(\Gamma))$ is torsion-free.*

Sketch of proofs of Proposition 5.27 and 5.28. In both cases, applying Corollary 5.16 at a fixed trivalent vertex v_0, the desired homology group A fits into a short exact sequence

$$0 \to \operatorname{coker} \delta_M \to A \to \ker \delta_{M-1} \to 0$$

for some M.

In the case of Proposition 5.27, we can use Proposition 5.25 and the explicit formula for the connecting homomorphism to show that the cokernel entry of the short exact sequence consists of the terms indexed by vertices other than v_0. Then by induction and an explicit calculation, the kernel entry splits and yields the term indexed by v_0.

In the case of Proposition 5.28, the kernel entry is torsion-free by induction so any torsion would have to come from the cokernel term. By explicitly examining two cases using the formula for the connecting homomorphism and Proposition 5.27 (the cases corresponding to whether the vertex indexing the summand in Proposition 5.27 is adjacent to v_0 or not), we conclude that the cokernel term is itself torsion-free. ☐

5.6 Case study: the complete graph K_4

We apply our results to give a complete description of $H_*(B(\mathsf{K}_4))$ as a $\mathbb{Z}[E]$-module, where K_4 is the complete graph on four vertices. As an intermediary result, we also give the corresponding computation for the net graph N.

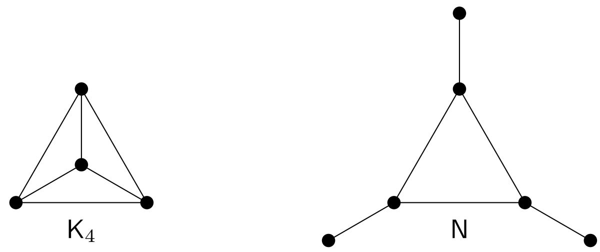

Figure 10: The complete graph K_4 and the net graph N

Notation 5.29. We number the vertices of K_4 and write e_{ab} for the edge connecting vertices a and b; thus, $e_{ab} = e_{ba}$ and e_{aa} is undefined. Write $h_b^{(a)}$ for the half-edge at a with edge e_{ab}. The action of the symmetric group Σ_4 on the vertices extends to an action on $S(\mathsf{K}_4)$ by bigraded chain maps intertwining the $\mathbb{Z}[E]$-action; indeed, a choice of parametrization defines an action of Σ_4 on K_4 by graph isomorphisms, and the Świątkowski complex is functorial. We write γ_a for the loop class avoiding vertex a and α_a for the star class at vertex a, with signs fixed as follows:

- $\gamma_4 = [h_2^{(1)} - h_3^{(1)} + h_3^{(2)} - h_1^{(2)} + h_1^{(3)} - h_2^{(3)}]$,
- $\alpha_4 = \alpha_{123}^{(4)} = [(h_1^{(4)} - h_2^{(4)})e_{34} + (h_3^{(4)} - h_1^{(4)})e_{24} + (h_2^{(4)} - h_3^{(4)})e_{14}]$,
- $\gamma_a = (-1)^{\mathrm{sgn}(\sigma)}\sigma(\gamma_4)$, and similarly for α_a, where σ is the cyclic permutation taking 4 to a.

We extend this notation to the net graph N by naming all 1-valent vertices 4; this is not ambiguous because we will never refer to the half-edges at these vertices and because every vertex is adjacent to at most one such vertex.

Proposition 5.30. *The homology $H_*(B(\mathsf{K}_4))$ is free Abelian and presented as a $\mathbb{Z}[E]$-module in terms of generators and relations as follows. Let $a, b, c, d \in \{1, 2, 3, 4\}$ be distinct.*

- *$H_0(B(\mathsf{K}_4))$ is generated by $\varnothing$ subject to relations identifying all edges.*
- *$H_1(B(\mathsf{K}_4))$ is generated by γ_a and α_a subject to the following relations:*

$$\sum_{a=1}^4 \gamma_a = 0 \qquad \gamma_a(e_{bc} - e_{cd}) = 0$$
$$\alpha_a = \alpha_b \qquad \gamma_a(e_{ab} - e_{ac}) = 0$$
$$\alpha_a(e - e') = 0 \qquad \gamma_a(e_{bc} - e_{ab}) = \alpha_b$$

- *$H_2(B(\mathsf{K}_4))$ is generated by $\gamma_a\alpha_a$ and $\alpha_a\alpha_b$ subject to the following relations:*

$$\begin{aligned} &\textstyle\sum_{a=1}^4 \gamma_a\alpha_a = 0 && \gamma_a\alpha_a(e_{bc} - e_{cd}) = 0 \\ &\alpha_a\alpha_b(e - e') = 0 \text{ if } e, e' \neq e_{ab} && \gamma_a\alpha_a(e_{bc} - e_{ab}) = \alpha_b\alpha_a \end{aligned}$$

- *$H_3(B(\mathsf{K}_4))$ is generated by $\alpha_a\alpha_b\alpha_c$ subject to the relations $\alpha_a\alpha_b\alpha_c(e_{ad} - e_{bd}) = 0$.*

- *$H_4(B(\mathsf{K}_4))$ is freely generated by $\alpha_1\alpha_2\alpha_3\alpha_4$.*

We will prove this result by relating K_4 to the net graph N, which is obtained by exploding a vertex. The corresponding result for N is the following.

LEMMA 5.31. *The homology $H_*(B(\mathsf{N}))$ is free Abelian and presented as a $\mathbb{Z}[E]$-module in terms of generators and relations as follows. Let $a, b, c \in \{1, 2, 3\}$ be distinct.*

- *$H_0(B(\mathsf{N}))$ is generated by $\varnothing$ subject to relations identifying all edges.*

- *$H_1(B(\mathsf{N}))$ is generated by γ_4 and α_a subject to the following relations:*

$$\begin{aligned} &\alpha_a(e - e') = 0 \text{ if } e, e' \neq e_{a4} && \gamma_4(e_{ab} - e_{ac}) = 0 \\ & && \gamma_4(e_{ab} - e_{a4}) = \alpha_a \end{aligned}$$

- *$H_2(B(\mathsf{N}))$ is generated by $\alpha_a\alpha_b$ subject to the relation $\alpha_a\alpha_b(e - e') = 0$ if $e, e' \notin \{e_{a4}, e_{b4}, e_{ab}\}$*

- *$H_3(B(\mathsf{K}_4))$ is freely generated by $\alpha_1\alpha_2\alpha_3$.*

Proof. This follows from Propositions 5.25, 5.27, and 5.6, along with the O, Q, and I-relations and a rank-counting argument (using Corollary 5.11, say). ☐

Proof of Proposition 5.30. The statement for H_0 holds for any connected graph, the statement for H_3 follows from Proposition 5.27, and the statement for H_4 follows from Proposition 5.25. For H_1, generation follows from Proposition 5.6 (along with the identification of loop classes containing four edges as the sum of two loop classes of the given form). The relation involving only loop classes follows from a relation in $H_1(\mathsf{K}_4)$ itself. The other stated relations follow directly from the Θ, I, Q, and O-relations. In fact, $H_1(B_n(\Gamma))$ is completely known in general [KP12]. The relevant part of that computation, which we reprove as Lemma C.9, shows by inspection of the relations that our presentation is complete. Thus, it remains to prove the statement for H_2. We may assume that $k \geq 2$, since H_2 vanishes in lower weight.

The outline of the proof is as follows. Using the long exact sequence of Corollary 5.16, we obtain an explicit $\mathbb{Z}$-linear description of H_2. Then, denoting by M the $\mathbb{Z}[E]$-module presented by generators and relations in the statement of the proposition, we exhibit a $\mathbb{Z}[E]$-linear map from M to H_2, which our integral description shows to be surjective. Finally, we verify that the two modules have the same integral rank in each weight.

The portion of the long exact sequence relevant to our purpose is

$$\cdots \to H_2(B_{k-1}(\mathsf{N}))^{\oplus 2} \xrightarrow{\delta_2} H_2(B_k(\mathsf{N})) \to H_2(B_k(\mathsf{K}_4)) \xrightarrow{\psi} \\ H_1(B_{k-1}(\mathsf{N}))^{\oplus 2} \xrightarrow{\delta_1} H_1(B_k(\mathsf{N})) \to \cdots$$

The kernel of δ_1 is free over $\mathbb{Z}$, so this portion of the sequence splits integrally, and we have $H_2(B_k(\mathsf{K}_4)) \cong \operatorname{coker}\delta_2 \oplus \ker\delta_1$ as $\mathbb{Z}$-modules. Passing to the cokernel of δ_2 merely identifies e_{a4} and e_{b4} in $H_2(B_k(\mathsf{N}))$ and thus contributes the space of external products $\alpha_a\alpha_b p(E)$ for $a, b \in \{1, 2, 3\}$, subject to the final relation $\alpha_a\alpha_b(e-e') = 0$ if $e, e' \neq e_{ab}$. On the other hand, $\ker\delta_1$ consists of the isomorphic image of the classes $e_{a4}^m e_{12}^{k-m-3}\gamma_a\alpha_a \in H_2(B_k(\mathsf{K}_4))$ for $a \in \{1, 2, 3\}$ and $0 \le m \le k-3$. Verifying this fact is a tedious calculation.

As for relations, the second, third, and fourth follow from the O-relation, the I-relation, and the Q-relation, respectively. For the first relation, counting ranks shows that $\sum_{a=1}^4 r_a\gamma_a\alpha_a = 0$ with at least one r_a non-zero, and symmetry under cyclic permutations implies that r_a is constant in a. The absence of torsion in $H_2(B(\mathsf{K}_4))$ forces r_a to be a unit.

For the rank calculation, we note that the $\operatorname{coker}\delta_2$ is the direct sum of three copies of the space of degree $k-4$ polynomials in two variables, which has integral rank $3(k-3)$, while $\ker\delta_1$ is the direct sum of three copies of the space of degree $k-3$ polynomials in two variables, which has integral rank $3(k-2)$. It follows that the integral rank of $H_2(B_k(\mathsf{K}_4))$ is $6k-15$, which is also a lower bound on the rank of M by surjectivity. Thus, it remains to show that the integral rank of M is at most $6k-15$. To see why this is the case, we note that the relations shown imply that M is spanned integrally in weight k by the classes $e_{ab}^{k-4-m}e^m\alpha_a\alpha_b$, where $1 \le a, b \le 4$, $0 \le m \le k-4$, and $e \neq e_{ab}$ is fixed, together with the classes $e_{bc}^{k-3}\gamma_a\alpha_a$, where $a, b, c \neq 4$. There are $6(k-3)$ generators of the former type and 3 of the latter type, so the rank is at most $6k-15$, as claimed. ☐

A Reminders on homotopy colimits

In this appendix, we present a summary of some relevant facts and definitions concerning homotopy colimits and related matters. For the general theory and full details, [Dug17], [GJ09] and [Hir02] are good references.

The objects of the simplicial indexing category Δ, being finite ordered sets, may naturally be regarded as categories, and the arrows of Δ determine functors among these categories.

Definition A.1. Let $\mathcal{D}$ be a category. The *nerve* of $\mathcal{D}$ is the simplicial set

$$N\mathcal{D}_\bullet := \operatorname{Fun}(\Delta^\bullet, \mathcal{D}).$$

A p-simplex $\sigma \in N\mathcal{D}_p$ is a string $\sigma(p) \to \cdots \to \sigma(0)$ of composable morpisms in $\mathcal{D}$. We say that $\mathcal{D}$ is *contractible* if the geometric realization $|N\mathcal{D}|$ is so.

Example A.2. A filtered or cofiltered category is contractible provided it is not empty. In particular, any category with an initial or terminal object is contractible.

Definition A.3. Let $F : \mathcal{D} \to \mathcal{T}\mathrm{op}$ be a functor. The *homotopy colimit* of F, denoted $\mathrm{hocolim}_{\mathcal{D}} F$, is the geometric realization of the simplicial space given in simplicial degree p by

$$\coprod_{\sigma \in N\mathcal{D}_p} F(\sigma(p))$$

with face and degeneracy maps induced by those of $N\mathcal{D}$.[1]

The homotopy colimit, thus defined, is functorial on the functor category, so that a natural transformation of functors $F \to G$ induces a map on homotopy colimits. A fundamental property of homotopy colimits is that they are *homotopy invariant*.

Lemma A.4. *The homotopy colimit of a natural weak homotopy equivalence between functors F and G is a weak homotopy equivalence.*

Definition A.5. Let X be a topological space and $\mathcal{U}$ an open cover of X. We say $\mathcal{U}$ is *complete* if it is possible to write any finite intersection of elements of $\mathcal{U}$ as a union of elements of $\mathcal{U}$.

We view the open cover $\mathcal{U}$ as partially ordered under inclusion and thereby as a category. There is a tautological functor $\Gamma : \mathcal{U} \to \mathcal{T}\mathrm{op}$ taking $U \in \mathcal{U}$ to the topological space U. Our main use of homotopy colimits is via the following result.

Theorem A.6 ([DI04, Prop. 4.6]). *If $\mathcal{U}$ is a complete cover of X, then the natural map*

$$\underset{\mathcal{U}}{\mathrm{hocolim}}\, \Gamma \longrightarrow X$$

is a weak homotopy equivalence.

One can also define homotopy colimits for functors valued in chain complexes. For simplicity we only define a special case; more detail is available in the references at the beginning of the appendix.

Given a simplicial chain complex $V : \Delta^{op} \to \mathcal{C}\mathrm{h}_{\mathbb{Z}}$, we may construct a bicomplex $\mathrm{Alt}(V)$ by taking the alternating sum of the face maps.

[1] Lemma A.4 implies that the homotopy colimit induces a functor from the *homotopy category* of functors from $\mathcal{D}$ to $\mathcal{T}\mathrm{op}$ to the homotopy category of $\mathcal{T}\mathrm{op}$. It is typical to either define the homotopy colimit as this latter functor or to define *a* homotopy colimit as any functor with similar properties to that in our definition which induces this same functor at the level of homotopy categories. We will not need this level of generality.

Proposition A.7 ([Dug17, Prop. 19.9]). *Let $V : \Delta^{op} \to \mathcal{C}h_{\mathbb{Z}}$ be a simplicial chain complex concentrated in non-negative degrees. There is a natural weak equivalence*

$$\operatorname*{hocolim}_{\Delta^{op}} V \simeq \mathrm{Tot}(\mathrm{Alt}(V)),$$

where Tot *denotes the total complex.*

The following standard result asserts that the notions of homotopy colimit valued in topological spaces and chain complexes are compatible.

Proposition A.8. *Let $F : \mathcal{D} \to \mathcal{T}op$ be a functor. There is a natural quasi-isomorphism*

$$\operatorname*{hocolim}_{\mathcal{D}} C^{\mathrm{sing}}(F) \simeq C^{\mathrm{sing}}(\operatorname*{hocolim}_{\mathcal{D}} F).$$

We shall also make use of a relative version of the homotopy colimit construction. Recall that, given a functor $T : \mathcal{D}_1 \to \mathcal{D}_2$ and the existence of enough colimits in $\mathcal{C}$, the restriction functor $T^* : \mathrm{Fun}(\mathcal{D}_2, \mathcal{C}) \to \mathrm{Fun}(\mathcal{D}_1, \mathcal{C})$ admits a left adjoint Lan, the *left Kan extension* functor. The homotopical version of this construction is the following.[2]

Definition A.9. Let $\mathcal{D}_2 \xleftarrow{T} \mathcal{D}_1 \xrightarrow{F} \mathcal{C}$ be functors with $\mathcal{C} = \mathcal{T}op$ or $\mathcal{C} = \mathcal{C}h_{\mathbb{Z}}$. The *homotopy left Kan extension* of F along T is the functor from $\mathcal{D}_2$ to $\mathcal{C}$ given (on objects) by

$$\mathrm{hoLan}_T\, F(d) = \operatorname*{hocolim}_{(T\downarrow d)}(F \circ \mathrm{forget}).$$

Here, for an object $d \in \mathcal{D}_2$, the *overcategory* $(T \downarrow d)$ has as objects the pairs (d', f) with $d' \in \mathcal{D}_1$ and $f : T(d') \to d$ a morphism in $\mathcal{D}_2$. A morphism from $f : T(d') \to d$ to $g : T(d'') \to d$ is a morphism $h : d' \to d''$ such that $g \circ T(h) = f$. The forgetful functor to $\mathcal{D}_1$ takes (d', f) to d'. The construction $(T \downarrow d)$ is functorial in d, and $\mathrm{hoLan}_T\, F$ extends to a functor using this functoriality and the functoriality of the homotopy colimit.

Example A.10. When $\mathcal{D}_2$ is the trivial category $\mathbf{1}$ with one object $*$ (so that there is a unique functor $* : \mathcal{D}_1 \to \mathbf{1}$), this construction recovers the homotopy colimit:

$$\mathrm{hoLan}_*\, F(*) \cong \operatorname*{hocolim}_{\mathcal{D}_1} F.$$

Dually, the objects of the *undercategory* $(d \downarrow T)$ has as objects the pairs (f, d') with $d' \in \mathcal{D}_1$ and $f : d \to T(d')$ a morphism in $\mathcal{D}_2$, and as morphisms morphisms h in $\mathcal{D}_1$ satisfying a dual condition. Overcategories and undercategories are very useful in the calculation of homotopy colimits. In order to say how this is so, we require a preliminary definition.

[2] As with the homotopy colimit, the homotopy left Kan extension may be defined more invariantly in terms of homotopy categories. With this setup, our definition is a proposition—see [Dug17, Prop. 10.2], for example.

Definition A.11. Let $T : \mathcal{D}_1 \to \mathcal{D}_2$ be a functor. We say that T is

1. *homotopy final* if $(d{\downarrow}T)$ is contractible for every $d \in \mathcal{D}_2$, or
2. *homotopy initial* if $(T{\downarrow}d)$ is contractible for every $d \in \mathcal{D}_2$.

That is, T is homotopy initial just in case $T^{op} : \mathcal{D}_1^{op} \to \mathcal{D}_2^{op}$ is homotopy final.

Proposition A.12 ([Rie14, Thm. 8.5.6]). *Let $\mathcal{D}_1 \xrightarrow{T} \mathcal{D}_2 \xrightarrow{F} \mathcal{C}$ be functors with $\mathcal{C} = \mathcal{T}\mathrm{op}$ (or $\mathcal{C} = \mathcal{C}\mathrm{h}_{\mathbb{Z}}$). If T is homotopy final, then the natural map*

$$\operatorname*{hocolim}_{\mathcal{D}_1} T^*F \to \operatorname*{hocolim}_{\mathcal{D}_2} F$$

is a weak homotopy equivalence (quasi-isomorphism).

We will make use of the following immediate consequences of this result.

Corollary A.13. *Let $\mathcal{D}$ be a category and $\underline{c} : \mathcal{D} \to \mathcal{C}$ the constant functor at $c \in \mathcal{C}$, where $\mathcal{C} = \mathcal{T}\mathrm{op}$ (or $\mathcal{C} = \mathcal{C}\mathrm{h}_{\mathbb{Z}}$). If $\mathcal{D}$ is contractible, then the natural map*

$$\operatorname*{hocolim}_{\mathcal{D}} \underline{c} \to c$$

is a weak homotopy equivalence (quasi-isomorphism).

Corollary A.14. *Let $T : \mathcal{D}_1 \to \mathcal{D}_2$ be any functor. If T is homotopy final, then the induced map $N\mathcal{D}_1 \to N\mathcal{D}_2$ is a weak homotopy equivalence.*

B Comparing decompositions

In this appendix, we tighten the combinatorial side of the correspondence between decompositions and relative tensor products arising from Theorem 3.20. We present a general conceptual framework in which comparison questions may be formulated, and we provide one uniform answer to such questions in the form of Proposition B.5. We then apply this comparison result in two examples used in the main text. We hope that these examples will make manifest that decompositions and relative tensor products may be manipulated in an essentially identical way.

B.1 Abstract setup

Throughout this section, we work exclusively with decompositions $\mathcal{E}$ and $\mathcal{F}$ of a fixed complex X. We assume that F is a symmetric monoidal functor determining both an $\mathcal{E}$-local invariant and an $\mathcal{F}$-local invariant. With more notational baggage, one could work in a more general setting.

Definition B.1. An (r, s)-*comparison scheme* is a functor $\alpha : \Delta^r \to \Delta^s$.

We think of a comparison scheme as specifying the pattern of an operation on iterated bimodules, as reflected in the simplicial bar construction, or, equally, the pattern of an operation on decompositions, as reflected in the trace.

NOTATION B.2. Given an r-fold decomposition $\mathcal{E}$, an s-fold decomposition $\mathcal{F}$, and an (r,s)-comparison scheme α, we write $\mathcal{G}_\alpha$ for the category of pairs (A,B) with $A \in \mathcal{G}_r$, $B \in \mathcal{G}_s$, $\gamma_{\mathcal{E}}(A) \subseteq \gamma_{\mathcal{E}}(B)$, and $\alpha(\tau(A)) = \tau(B)$. By design, this category fits into a commuting diagram

$$\begin{array}{ccccc} \mathcal{G}_r & \xleftarrow{\;p\;} & \mathcal{G}_\alpha & \xrightarrow{\;q\;} & \mathcal{G}_s \\ {\scriptstyle\tau}\downarrow & & & & \downarrow{\scriptstyle\tau} \\ \Delta^r & & \xrightarrow{\quad\alpha\quad} & & \Delta^s. \end{array}$$

DEFINITION B.3. We say that $\mathcal{E}$ and $\mathcal{F}$ are *α-comparable* if $p : \mathcal{G}_\alpha \to \mathcal{G}_r$ is homotopy initial.

Roughly, α-comparability asserts that every subspace in the image of $\gamma_{\mathcal{E}}$ is contained in a contractible collection of subspaces in the image of $\gamma_{\mathcal{F}}$ with compatible patterns of connected components.
We now add local flows into the mix.

DEFINITION B.4. Suppose that F is equipped with isotopy invariant local flows on both $\mathcal{E}$ and $\mathcal{F}$.

1. We say that a comparison scheme α is *flow compatible* if, for every $(A,B) \in \mathcal{G}_\alpha$, the inclusion $\gamma_{\mathcal{E}}(A) \subseteq \gamma_{\mathcal{F}}(B)$ is flow compatible.

2. Fix a flow compatible comparison scheme α. A *comparison datum* is a map $\mathrm{Bar}_\Delta(I(\mathcal{E})) \to \mathrm{Bar}_\Delta(I(\mathcal{F})) \circ \alpha^{op}$ in the homotopy category of r-fold simplicial chain complexes, which fits into the commuting diagram

$$\begin{array}{ccc} p^*I(F(\gamma_{\mathcal{E}})) & \longrightarrow & q^*I(F(\gamma_{\mathcal{F}})) \\ & & \downarrow{\scriptstyle q^*\psi} \\ {\scriptstyle p^*\psi}\downarrow & & q^*(\mathrm{Bar}_\Delta(I(\mathcal{F})) \circ \tau^{op}) \\ & & \| \\ p^*(\mathrm{Bar}_\Delta(I(\mathcal{E})) \circ \tau^{op}) & \dashrightarrow & p^*(\mathrm{Bar}_\Delta(I(\mathcal{F})) \circ \alpha^{op} \circ \tau^{op}). \end{array}$$

The upper horizontal arrow exists by flow compatibility. We call the map $\mathrm{Bar}(I(\mathcal{E})) \to \mathrm{Bar}(I(\mathcal{F}))$ induced by a comparison datum the "comparison map." We these definitions in hand, the following result is essentially formal.

PROPOSITION B.5. *If $\mathcal{E}$ and $\mathcal{F}$ are α-comparable, then the composite*

$$\mathrm{Bar}(I(\mathcal{E})) \to \mathrm{Bar}(I(\mathcal{F})) \simeq C^{\mathrm{SD}}(F(X))$$

of the comparison map followed by the weak equivalence of Theorem 3.20 applied to $\mathcal{F}$ coincides with the weak equivalence of Theorem 3.20 applied to $\mathcal{E}$. In particular, the comparison map is a quasi-isomorphism.

B.2 Parenthesization

For our first sample application, we begin by noting that the derived tensor product computed by the bar construction on an iterated bimodule enjoys a great deal of symmetry, remaining invariant under any parenthesization of the factors. To give one example, there is a canonical quasi-isomorphism

$$M_0 \otimes^{\mathbb{L}}_{R_1} M_1 \otimes^{\mathbb{L}}_{R_2} M_2 \simeq M_0 \otimes^{\mathbb{L}}_{R_1} (M_1 \otimes^{\mathbb{L}}_{R_2} M_2).$$

The corresponding comparison in the context of the decomposition theorem is beween, on the one hand, the decomposition $\mathcal{E}$ with components X_0, X_1, and X_2 and bridges $A_1 \times I$ and $A_2 \times I$; and, on the other hand, the decomposition $\mathcal{E}_{\mathrm{par}}$ with components X_0 and $X_1 \amalg_{A_2\times I} X_2$ and bridge $A_1 \times I$ (see Figure 11). In order to apply Proposition B.5, we proceed as follows.

1. We choose our $(2,1)$-comparison scheme $\alpha : \Delta^2 \to \Delta$ to be the projection onto the first factor.
2. With this choice, $\mathcal{E}$ and $\mathcal{E}_{\mathrm{par}}$ are α-comparable. Indeed, the undercategory in question is filtered.
3. We assume we are given isotopy invariant local flows for which α is flow compatible. This point will depend on the specifics of the situation.
4. In the manner of Construction 3.16, we build a collection of maps

   $$I(F(X_1)) \otimes I(F(A_2)) \otimes \cdots \otimes I(F(A_2)) \otimes I(F(X_2)) \to I(F(X_1 \amalg_{A_2\times I} X_2))$$

   using isotopy invariance and flow compatibility of the structure maps of F and the comparison scheme. These maps respect the simplicial structure maps, and the induced map of simplicial chain complexes is a comparison datum.

Invoking Proposition B.5, we see that the quasi-isomorphism

$$\mathrm{Bar}(I(\mathcal{E})) \xrightarrow{\sim} \mathrm{Bar}(I(\mathcal{E}_{\mathrm{par}}))$$

is compatible with those supplied by Theorem 3.20.

Remark B.6. It should be clear that an arbitrary parenthesization may be treated in an identical fashion.

B.3 Folding

Our second example concerns the canonical quasi-isomorphism

$$M_0 \otimes^{\mathbb{L}}_{R_1} M_1 \otimes^{\mathbb{L}}_{R_2} M_2 \simeq (M_0 \otimes M_2) \otimes^{\mathbb{L}}_{R_1 \otimes R_2^{op}} M_1$$

The corresponding comparison in the context of the decomposition theorem is between the decomposition $\mathcal{E}$ already introduced and the decomposition $\mathcal{E}_{\mathrm{fld}}$ with components $X_0 \amalg X_2$ and X_1 and bridge $(A_1 \amalg A_2) \times I$ obtained by "folding the end over" (see Figure 11). The data in the this case is the following.

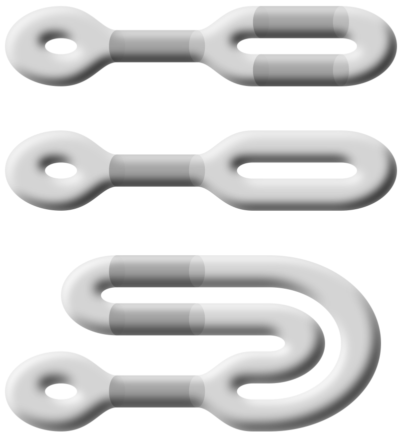

Figure 11: From top to bottom, the decompositions $\mathcal{E}$, $\mathcal{E}_{\text{par}}$, and $\mathcal{E}_{\text{fld}}$

1. We choose our $(1,2)$-comparison scheme $\alpha : \Delta \to \Delta^2$ to be the identity in the first factor and the "flip" automorphism in the second.
2. With this choice, $\mathcal{E}_{\text{fld}}$ and $\mathcal{E}$ are α-comparable. Indeed, the undercategory in question has a final object.
3. We assume we are given isotopy invariant local flows for which α is flow compatible. This point will depend on the specifics of the situation.
4. Observe that the simplicial chain complexes $\text{Bar}_\Delta(\mathcal{E}_{\text{fld}})$ and $\text{Bar}_\Delta(\mathcal{E})\circ\alpha^{op}$ are identical in each simplicial degree, except that the former uses the local flow for $\mathcal{E}_{\text{fld}}$, while the latter uses the local flow for $\mathcal{E}$. The comparison datum comes from matching terms and using the flow compatibility of α.

Invoking Proposition B.5, we see that the quasi-isomorphism

$$\text{Bar}(I(\mathcal{E}_{\text{fld}})) \xrightarrow{\sim} \text{Bar}(I(\mathcal{E}))$$

is compatible with those supplied by Theorem 3.20.

C Degree one homology of graph braid groups

The homology group $H_1(B_k(\Gamma))$ for a general connected graph Γ is completely understood; indeed, according to a theorem of Ko–Park [KP12, Thm. 3.16], this group may be identified solely in terms of connectivity and planarity data from Γ. Their argument proceeds through an intricate combinatorial and linear algebraic analysis of the discrete Morse data constructed in [FS05], which quickly becomes very technical.

In this appendix, we use the Świątkowski complex to simplify the calculation of H_1 by giving streamlined proofs of four key lemmas of [KP12], which appear in §C.3 below, cross-referenced with the corresponding results in the original text (and in [HKRS14], where versions of them also appear). These four results directly imply the calculation of H_1, which we outline below but do not state in full. We make no claim of originality in this appendix, our purpose being only to demonstrate the efficiency of the Świątkowski complex in applications.

C.1 Cuts and connectivity

In order to proceed, we will require some terminology from graph theory. Note that these invariants should only be used as defined on *simple* graphs (see §4.1) and may behave unexpectedly on general graphs.

Definition C.1. Let Γ be a simple graph. A *k-cut* is a set of k vertices whose removal topologically separates at least two vertices of Γ. A simple graph is *k-connected* if it has at least $k+1$ vertices and no $(k-1)$-cuts.
Given a 1-cut v in Γ, a *v-component* of Γ is the closure in Γ of a connected component of the complement of v in Γ.

The importance of connectivity for our purposes arises from the following classical result, called Menger's theorem.

Theorem C.2 ([Men27]). *Let $k > 0$. A simple graph is k-connected if and only if, for distinct vertices x and y in Γ, there exist k paths from x to y in Γ, disjoint except at endpoints.*

In other words, a simple graph is k-connected (for positive k) if and only if any pair of vertices x and y there is a graph embedding of a subdivision of the theta graph Θ_k into Γ taking the two vertices of the theta graph to x and y.
As the following result shows, high connectivity places strong constraints on the behavior of the first homology.

Proposition C.3. *If Γ is simple and 3-connected, then any two star classes in $H_*(B(\Gamma))$ coincide up to sign.*

Proof. By 3-connectivity, any two vertices of Γ are joined by three distinct paths, so a star class at one vertex coincides with some star class at any other vertex by the Θ-relation. Thus, it suffices to verify that two star classes α and α' at a fixed vertex v coincide up to sign. We may assume further that α and α' are induced by inclusions of S_3 differing only at a single half-edge, so that we have an inclusion of S_4 into Γ (with endpoint vertices $v_1, \ldots, v_4$ and central vertex v all distinct by 3-connectedness). We will show that the four star classes in Γ obtained in this way agree up to sign.
A corollary of Menger's theorem is the fact that in a k-connected graph, any two sets of vertices, each of size k, can be joined by k disjoint paths [BM08, Prop. 9.4]. Since Γ_v is 2-connected, two disjoint paths join $\{v_1, v_4\}$ and $\{v_2, v_3\}$. By relabelling suppose one joins v_1 to v_2 and the other joins v_3 to v_4. Similarly,

two disjoint paths join $\{v_1, v_2\}$ to $\{v_3, v_4\}$. By switching the labels on v_1 and v_2 if need be, we may assume one joins v_1 to v_3 and the other joins v_2 to v_4. We now have two extensions of the inclusion $\mathsf{S}_4 \to \Gamma$ to an inclusion of a subdivision of the figure-eight graph $\mathsf{8}$ as in Figure 12, and the Q-relation and I-relation imply that $\alpha_{123} = \alpha_{124}$ and that $\alpha_{432} = \alpha_{431}$ (using the first set of paths), and that $\alpha_{312} = \alpha_{314}$ (using the second set of paths). □

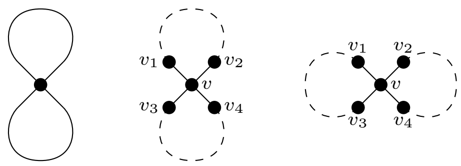

Figure 12: The figure-eight graph $\mathsf{8}$ and two inclusions of subdivisions of it at a vertex of valence at least 4 in a triconnected graph

C.2 Minors and planarity

Recall that a graph is said to be *planar* if it can be embedded in $\mathbb{R}^2$. The goal of this section is to establish the following connection between planarity and configuration spaces.

Lemma C.4. *Let Γ be a 3-connected simple graph. If Γ is non-planar, then any star class of $H_1(B_2(\Gamma))$ is 2-torsion.*

As before, the proof will proceed by reduction to atomic cases. In order to see how this reduction will proceed, we require an auxuiliary notion.

Definition C.5. Let Γ and Γ' be graphs. We say that Γ is a *minor* of Γ' if Γ may be obtained up to isomorphism from Γ by repeated application of the following operations:

1. contract an edge;
2. remove an edge;
3. remove an isolated vertex.

The relevance of minors for our purposes is the following classical result (see Figures 13 and 14 for terminology).

Theorem C.6 ([Wag37]). *A graph is non-planar if and only if it admits K_5 or $\mathsf{K}_{3,3}$ as a minor.*

In order to apply this criterion, we must first clarify the relationship between the Świątkowski complex of a graph and that of its minors. We begin by noting that, if Γ is a minor of Γ', then E is naturally identifed with a subset of E'.

Lemma C.7. *Let Γ be a minor of Γ'. There is a map $f : S(\Gamma) \to S(\Gamma')$ of differential graded $\mathbb{Z}[E]$-modules such that $f_*(\alpha)$ is a star class whenever α is.*

Proof. We may assume that Γ is obtained from Γ' by a single application of one of the operations of Definition C.5. In the case of operation (2) or (3), Γ is a subgraph of Γ', and the claim is immediate from the standard functoriality of the Świątkowski complex. Since the same holds for the contraction of a tail or a self-loop, we may assume that Γ is obtained from Γ' by contracting an edge with two distinct vertices v_1 and v_2. Denote the corresponding half-edges by h_1 and h_2, and let v_0 be the vertex that is the image of the closure of e under the quotient map $\Gamma' \to \Gamma$.

In order to define f, we note that, under the quotient map, each vertex $v \neq v_0$ of Γ is canonically identified with a vertex $\tilde{v}$ of Γ', and each half-edge of h of Γ is canonically identified with a half-edge $\tilde{h}$ of Γ'. With this in mind, we set

$$f(v) = \begin{cases} \tilde{v}, & v \neq v_0 \\ e, & v = v_0 \end{cases} \qquad f(h) = \begin{cases} \tilde{h}, & v(h) \neq v_0 \\ \tilde{h} - h_j, & v(\tilde{h}) = v_j. \end{cases}$$

By inspection, f respects the differential and $\mathbb{Z}[E]$-action, and the claim regarding star classes is a direct calculation with star cycles. □

Proof of Lemma C.4. Since Γ is non-planar, it admits K_5 or $\mathsf{K}_{3,3}$ as a minor, and we may push a star class of the minor forward using Lemma C.7 to obtain a star class in Γ. Therefore, it suffices to prove the claim for K_5 and $\mathsf{K}_{3,3}$.

In the case of K_5, we can find an embedded copy of the theta graph Θ_3 as in Figure 13. This shows that the star class α_{123} at v is the negative of the star class $\alpha_{3'1'2'} = \alpha_{1'2'3'}$ at v'. The picture depicted and concomitant Θ-relation can be rotated by $\frac{2\pi}{5}$; applying such relations 5 times shows that the star class α_{123} is equal to $(-1)^5\alpha_{123}$.

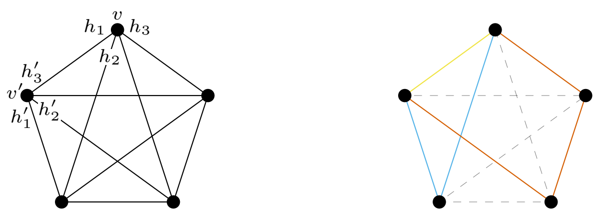

Figure 13: The complete graph K_5 and an embedded subdivision of Θ_3 with edges of Θ_3 color-coded.

In the case of $\mathsf{K}_{3,3}$, we have two embedded subdivisions of Θ_3, as in Figure 14. Applying the Θ-relation twice, we see that both a star class at the upper left vertex and its additive inverse are equal to the same star class at the upper right vertex. □

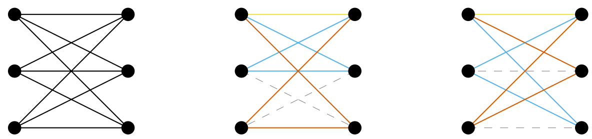

Figure 14: The complete bipartite graph $\mathsf{K}_{3,3}$ and two embedded subdivisions of Θ_3 with edges of Θ_3 color-coded.

C.3 Four lemmas of Ko–Park

We are now equipped to prove the results of [KP12] alluded to above, which together easily yield a complete calculation of $H_1(B_k(\Gamma))$ for an arbitrary connected graph Γ. We content ourselves with an outline of this computation, directing the interested reader to [KP12] for details. Versions of these lemmas are also proven in [HKRS14], and we have included citations to the appropriate parts of that paper as well. Harrison et al. use combinatorial ingredients similar to ours, but the logic of their arguments often diverges substantially from ours.

Since any graph has a simple subdivision (see §4.1), we may assume that Γ is simple. In this case, there are classical decomposition theorems from graph theory associating to a 1-connected simple graph a set of 2-connected simple graphs, obtained by repeated 1-cuts, and to each of these a set of 3-connected simple graphs, obtained by repeated 2-cuts. With this tool in hand, together with the observation that it suffices to assume that Γ is 1-connected, the argument proceeds as follows.

1. In characterizing the effect of a 1-cut, Lemma C.10 reduces the computation for Γ to the computation for each associated 2-connected graph.

2. Assuming now that Γ is 2-connected, Lemma C.8 grants that $H_1(B_k(\Gamma))$ is independent of k for $k \geq 2$.

3. Lemma C.14 describes the effect of a 2-cut on $H_1(B_2(\Gamma))$ for 2-connected Γ, reducing the calculation to graphs that are simpler in some technical sense. A variant of a classical decomposition theorem in graph theory [Tut66, CE80] says iterating this simplification procedure eventually reduces all 2-connected graphs to 3-connected graphs, cycle graphs, and theta graphs—see Remark C.12.

4. Finally, Lemma C.9 computes $H_1(B_2(\Gamma))$ for 3-connected Γ. The calculation for theta graphs can be understood by inspection or via the vertex long exact sequence of Corollary 5.16.

With careful bookkeeping of cuts and these graph decomposition theorems, one can assemble these results to arrive at an explicit answer, which we do not repeat here.

The first result is a sharpened version of the degree 1 component of Proposition 5.21 in the presence of sufficient connectivity.

Lemma C.8 ([KP12, Lem. 3.12]; see also [HKRS14, Theorem 5]). *Suppose that* Γ *is simple and* 2*-connected. For any* $e \in E$ *and* $k \geq 2$*, multiplication by* e^{k-2} *induces an isomorphism*

$$H_1(B_2(\Gamma)) \xrightarrow{\simeq} H_1(B_k(\Gamma)).$$

Proof. By Propositions 5.21 and 5.6, it will suffice to show that $p(E)\alpha$ and $p(E)\gamma$ are divisible by e in the expected range, where $p(E)$ is a monomial in E, α is a star class at the vertex v, and γ is a loop class.

In the case of a star class α, there is nothing to show provided $p(E) \in \mathbb{Z}[e]$, so we may write $p(E)\alpha = [q(E)e'a]$, where $e' \neq e$ and $[a] = \alpha$. Since Γ is 2-connected, both e' and e have vertices distinct from v, and there is an injective path in $\Gamma \setminus \{v\}$ between them. Letting h denote the alternating sum of the half-edges involved in this path, we have $\partial(q(E)ha) = q(E)ea - q(E)e'a$. An easy induction on the degree of q completes the argument in this case.

In the case of a loop class γ, we write $p(E)\gamma = [q(E)e'c]$, where $e' \neq e$ and $[c] = \gamma$. Since Γ is connected, there is a path connecting e' to e. If this path contains none of the edges or vertices involved in c, then the same argument as before shows that $q(E)ec \sim q(E)e'c$; on the other hand, if the path is contained entirely in c, then the same conclusion follows by the O-relation. Thus, we may assume that e lies in the complement of c, e' lies in c, and that the two share a vertex. By the Q-relation, we have $q(E)ec \sim \pm q(E)e'c \pm q(E)a$, where $[a]$ is a star class. Since the case of a star class is known, this completes the proof. ☐

The second result is the base calculation, the 3-connected case.

Lemma C.9 ([KP12, Lem. 3.15]; see also [HKRS14, Lemma 3, Theorem 4]). *If* Γ *is* 3*-connected, then* $H_1(B_2(\Gamma)) \cong \mathbb{Z}^{\beta_1(\Gamma)} \oplus K$*, where*

$$K = \begin{cases} \mathbb{Z} & \text{if } \Gamma \text{ is planar} \\ \mathbb{Z}/2 & \text{else.} \end{cases}$$

Here K is generated by a star class and $\mathbb{Z}^{\beta_1(\Gamma)}$ by the inclusion of cycle graphs into Γ.

Proof. Choose a set of loops c_i in Γ representing a basis for $H_1(\Gamma)$ and a star class α. Define a map

$$\left(\bigoplus_{i=1}^{\beta_1(\Gamma)} H_1(B_2(c_i)) \right) \oplus K \cong \mathbb{Z}^{\beta_1(\Gamma)} \oplus K \xrightarrow{\psi} H_1(B_2(\Gamma))$$

using the inclusions of c_i in Γ and taking the generator of K to α. This map is well-defined by Lemma C.4 and surjective by Propositions 5.6 and C.3.

Next, we "thicken" Γ, replacing vertices with disks and edges with strips, obtaining a surface Σ of genus g with b boundary components equipped with an embedding $\iota : \Gamma \to \Sigma$ that is also a homotopy equivalence (see [Kon92, p. 4-5], for example, and the references therein for more on this classical tactic). Note that $g > 0$ if Γ is non-planar, while we may take $g = 0$ if Γ is planar. Moreover, by comparing Euler characteristics, we find that

$$\beta_1(\Gamma) = 2g + b - 1.$$

Now, from the explicit presentation for $\pi_1(B_2(\Sigma))$ given in [Bel04], it is easy to see that ι induces a surjection $\pi_1(B_2(\Gamma)) \to \pi_1(B_2(\Sigma))$ on fundamental groups, and thus also on first homology. By direct computation, the same presentation gives

$$\begin{aligned} H_1(B_2(\Sigma)) &\cong \begin{cases} \mathbb{Z} \oplus \mathbb{Z}^{b-1} & g = 0 \\ \mathbb{Z}/2 \oplus \mathbb{Z}^{2g+b-1} & g \geq 1 \end{cases} \\ &\cong \mathbb{Z}^{\beta_1(\Gamma)} \oplus K, \end{aligned}$$

Thus, $H_1(B_2(\Gamma))$ both admits a surjection from and a surjection onto $\mathbb{Z}^{\beta_1(\Gamma)} \oplus K$, which implies the claim by finite dimensionality. □

The third result describes the effect on H_1 of a 1-cut.

Lemma C.10 ([KP12, Lem. 3.11]; see also [HKRS14, Section 4.4 and Section 7]). *Let Γ be a graph, v a 1-cut of valence ν in Γ, and $\{\Gamma^{(i)}\}_{i=1}^{\mu}$ the set of v-components of Γ_v. There is an isomorphism*

$$H_1(B_k(\Gamma)) \cong \left(\bigoplus_{i=1}^{\mu} H_1(B_k(\Gamma^{(i)})) \right) \oplus \mathbb{Z}^{N(k,\Gamma,v)}$$

where

$$N(k, \Gamma, v) = (\nu - 2)\binom{k+\mu-2}{k-1} - \binom{k+\mu-2}{k} - (\nu - \mu - 1).$$

Proof. Applying Corollary 5.16 at v, we obtain the exact sequence

$$\begin{aligned} \cdots \to \bigoplus^{|H(v)|-1} H_1(B_{k-1}(\Gamma_v)) \xrightarrow{\delta_1} H_1(B_k(\Gamma_v)) \to \\ H_1(B_k(\Gamma)) \to \bigoplus^{|H(v)|-1} H_0(B_{k-1}(\Gamma_v)) \xrightarrow{\delta_0} H_0(B_k(\Gamma_v)) \to \cdots \end{aligned}$$

Since zeroth homology is free, $\ker \delta_0$ is as well, so $H_1(B_k(\Gamma)) \cong \operatorname{coker} \delta_1 \oplus \ker \delta_0$. Assume Γ is connected for simplicity; then $\pi_0(B_k(\Gamma_v))$ is in bijection with the set of partitions of k into μ distinguished blocks and we conclude by exactness that

$$\operatorname{rk}\ker\delta_0 = (\nu-1)\binom{k+\mu-2}{k-1} - \binom{k+\mu-1}{k} + 1.$$

A similar argument shows $H_1(B_k(\Gamma^{(i)})) \cong \operatorname{coker}\delta_1^{(i)} \oplus \ker\delta_0^{(i)}$ and $\bigoplus_i \ker\delta_0^{(i)} \cong \mathbb{Z}^{(\nu-\mu)}$. Then, after a little combinatorial rearrangement, all that remains is to show that $\operatorname{coker}\delta_1 \cong \bigoplus_{i=1}^{\mu}\operatorname{coker}\delta_1^{(i)}$. The maps $H_1(B_k(\Gamma_v^{(i)})) \to H_1(B_k(\Gamma_v))$ arising from the various inclusions induce a map $\bigoplus_{i=1}^{\mu}\operatorname{coker}\delta_1^{(i)} \to \operatorname{coker}\delta_1$. We will show that this map is an isomorphism.

By the Künneth formula (again using the fact that zeroth homology is free) the homology $H_1(B_k(\Gamma_v))$ splits as a direct sum over i. The ith summand consists of homology classes in $H_1(B_{k'}(\Gamma_v^{(i)}))$ equipped with an ordered partition of $k-k'$ into $\mu-1$ blocks for some $k' \leq k$. Passing to the quotient $\operatorname{coker}\delta_1$, every equivalence class in the ith summand has a representative homology class with $k' = k$. Since this representative is in the image of the map in question, surjectivity follows.

For injectivity, we note that the map $H_1(B_k(\Gamma_v^{(i)})) \to H_1(B_k(\Gamma_v))$ lands in the ith summand of the direct sum decomposition, and δ_1 respects this decomposition. Now for fixed i, we change basis for the direct sum indexing the domain of δ_1. The given basis is spanned by all half-edges of v except a special half-edge h_0. We choose the special half-edge to lie in the v-component $\Gamma_v^{(i)}$. Then for each v-component $\Gamma_v^{(j)}$ other that $\Gamma_v^{(i)}$, we keep one basis half-edge h_{j0} corresponding to an edge in $\Gamma_v^{(j)}$ and change basis to $h_{jk} - h_{j0}$ for each other basis half-edge h_{jk} in that v-component. We retain all basis half-edges in the v-component $\Gamma_v^{(i)}$ without change. All of the changed basis elements are in the kernel of δ_1. The retained basis element h_{j0} from other v-components each change one of the entries of the ordered partition of $k-k'$ under δ_1. The set of h_{j0} and the set of $\mu-1$ blocks are in bijection under this correspondence. A homology class α in the codomain $H_1(B_k(\Gamma_v))$ of δ_1 which is in the image of the inclusion of $H_1(B_k(\Gamma_v^{(i)}))$ has $k' = k$. Suppose α is in the image of δ_1, say $\alpha = \delta_1(p_j h_{j0} + q)$, with q in the summands indexed by half-edges in $\Gamma_v^{(i)}$. Then since $k' = k$ each p_j is zero, so that α is also in the image of $\delta_1^{(i)}$. □

The final result describes the effect on H_1 of a 2-cut. We will establish a result for a certain type of decomposition.

Notation C.11. Let $\{x,y\}$ be a 2-cut in a 2-connected simple graph Γ. An $\{x,y\}$-*decomposition* consists of a (redundant) collection of subgraphs of Γ, namely $(\mathsf{G}_1, \mathsf{G}_{-1}, \mathsf{P}_1, \mathsf{P}_{-1}, \Gamma_1, \Gamma_{-1}, \mathsf{C})$, where

1. the subgraph G_1 contains x and y and $\mathsf{G}_1\backslash\{x,y\}$ is a connected component of $\Gamma\backslash\{x,y\}$,
2. the subgraph G_{-1} is $(\Gamma\backslash\mathsf{G}_1)\cup\{x,y\}$,

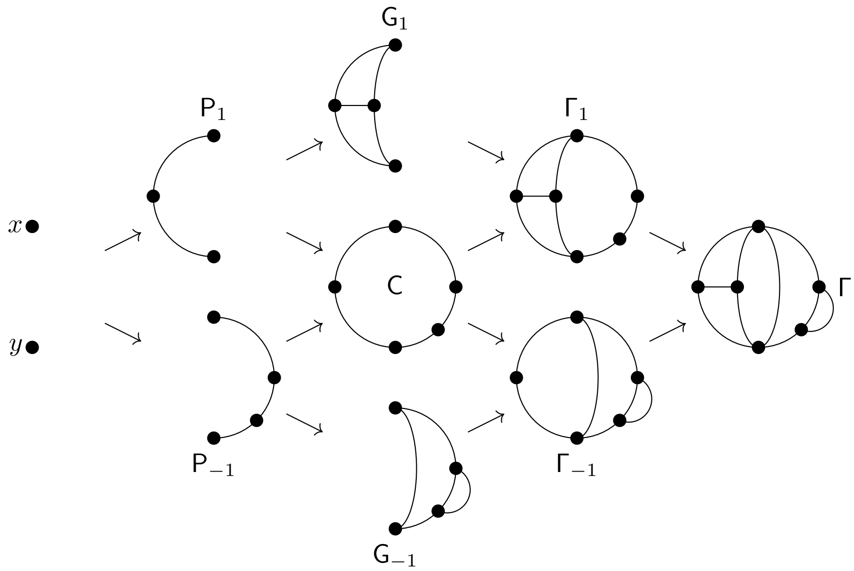

Figure 15: An example of an $\{x, y\}$-decomposition. Every diamond is a pushout in $\mathcal{T}op$.

3. the subgraph P_i is a simple path in G_i between x and y,
4. the subgraph Γ_i is the union of G_i and P_{-i},
5. the subgraph C is the union of P_1 and P_{-1}.

See Figure 15 for an example. This notation is not symmetric in its indices because of the connectivity requirement on $\mathsf{G}_1 \backslash \{x, y\}$. Note that the graphs $\Gamma_{\pm 1}$ are still 2-connected and simple.

Remark C.12. This type of decomposition comes from classical graph theory [Tut66, CE80], although our notation and terminology differ. In particular, we use a different definition of 2-cut which coincides with the classical definition on simple graphs, and our $\Gamma_{\pm 1}$ are subdivisions of the components considered in the above references.
Using the argument of [CE80, Theorem 1], for example, we can make a sequence of carefully chosen $\{x, y\}$-decompositions to decompose Γ into a collection of subdivisions of cycle graphs, θ-graphs, and 3-connected graphs. Thus, since cycles and θ-graphs may be treated by ad hoc means, reducing the problem of understanding $H_1(B(\Gamma))$ to that of understanding $H_1(B(\Gamma_{\pm 1}))$, as Lemma C.14 below, is tantamount to reducing the problem of understanding $H_1(B(\Gamma))$ for 2-connected Γ to the problem of understanding $H_1(B(\Gamma'))$ for 3-connected Γ'. See [KP12] for details of this argument.

LEMMA C.13. *Let $\mathsf{\Gamma}$ be a 2-connected simple graph with a 2-cut $\{x, y\}$ and an $\{x, y\}$-decomposition. The natural map of Świątkowski complexes $S(\mathsf{P}_i) \to S(\mathsf{G}_i)$ admits a retraction $S(\mathsf{G}_i) \to S(\mathsf{P}_i)$ with the property that generators of $S(\mathsf{G}_i)$ which avoid x and its half-edges (respectively y and its half-edges) are taken to sums of generators of $S(\mathsf{P}_i)$ which avoid x and its half-edge (respectively y and its half-edge).*

We think of $\mathsf{G}_{\pm 1}$ and $\mathsf{P}_{\pm 1}$ as relatively simpler subgraphs of the $\{x, y\}$-decomposition. In Lemma C.14, we will use these retractions to build similar retractions involving the more complicated subgraphs $\mathsf{\Gamma}$ and $\mathsf{\Gamma}_{\pm 1}$.

Proof. Suppose we are given a retraction $f : S_{\leq m}(\mathsf{G}_i) + S(\mathsf{P}_i) \to S(\mathsf{P}_i)$ with the desired property. The problem of extending this map to a generator b of degree $m + 1$ is that of choosing a nullhomotopy in $S(\mathsf{P}_i)$ for $f(\partial(b))$ having the desired property. Since $S(\mathsf{P}_i)$ and its reductions at x, at y, and at $\{x, y\}$ are all contractible in each weight, this extension is unobstructed as long as $m > 0$. Thus, by induction on m, it will suffice to produce a retraction $f : S_{\leq 1}(\mathsf{G}_i) + S(\mathsf{P}_i) \to S(\mathsf{P}_i)$ with the desired property. For this, we extend the identity on $S_{\leq 1}(\mathsf{P}_i)$ by

1. sending each edge or vertex not lying in P_i to a fixed arbitrarily chosen edge e of P_i,
2. sending each half-edge whose vertex is not in P_i to zero, and
3. sending each half-edge whose vertex v is in P_i to a chain in $S(\mathsf{P}_i)$ (avoiding x and y except potentially at v) whose boundary is the difference between v and e. This is possible because P_i is topologically an interval with endpoints x and y. □

LEMMA C.14 ([KP12, Lem. 3.13]; see also [HKRS14, Section 4.3]). *Let $\mathsf{\Gamma}$ be a 2-connected simple graph with a 2-cut $\{x, y\}$ and an $\{x, y\}$-decomposition. Then the sequence*

$$0 \to H_1(B_2(\mathsf{C})) \to H_1(B_2(\mathsf{\Gamma}_1)) \oplus H_1(B_2(\mathsf{\Gamma}_{-1})) \to H_1(B_2(\mathsf{\Gamma})) \to 0$$

is split exact, where the maps are induced by the respective inclusions, with the map $H_1(B_2(\mathsf{\Gamma}_{-1})) \to H_1(B_2(\mathsf{\Gamma}))$ twisted by an overall sign.

Proof. The composite of the middle two maps is the sum of the map induced by the inclusion of C into $\mathsf{\Gamma}$ with its negative, hence zero. For injectivity of the second map, we appeal to the argument of Lemma C.9, which shows that the loop class corresponding to C is non-zero. For surjectivity of the third map, by 2-connectedness of $\mathsf{\Gamma}$, any graph morphism of S_3 to $\mathsf{\Gamma}$ can be extended to an embedding of some lollipop graph. By Proposition 5.6 and the Q-relation it suffices to show that classes of the form $e\gamma$ lie in the image, where e is an edge and γ a loop class.

If e does not lie in the representing loop of γ, then by 2-connectedness of Γ we may write $e\gamma = e\gamma' + e\gamma''$ where e lies within the representing loops of both γ' and γ''. By the O-relation, any two choices for e within γ' represent the same class. This reduces us to classes $e\gamma$ where γ passes through both G_1 and G_{-1} and e lies in G_1. Then $e\gamma$ can be rewritten as $e\gamma_1 + e\gamma_{-1}$, where γ_i is a loop class in Γ_i for $i \in \{-1, 1\}$. The first of these lies in the image of $H_1(B_2(\Gamma_1))$, and the second lies in the image of $H_1(B_2(\Gamma_{-1}))$ by the connectivity property of G_1.

It remains to show exactness and splitting, for which we will use the retracts of Lemma C.13. A generator of $S(\Gamma)$ is a product of edges, half-edges, and vertices of G_1 and G_{-1} (which intersect only at the vertices x and y), and so, by applying these retracts on each of these two pieces, we obtain retracts $\pi_i : S(\Gamma) \to S(\Gamma_i)$ and $\rho_i : S(\Gamma_i) \to S(\mathsf{C})$ of the maps induced by the respective inclusions. We only have well-definition of these retracts because of the "boundary conditions" on generators involving x and y in the Lemma C.13. Moreover, because the maps of Lemma C.13 are *retracts*, the following diagram commutes:

$$\begin{array}{ccc} S(\Gamma_i) & \xrightarrow{\rho_i} & S(\mathsf{C}) \\ \downarrow & & \downarrow \\ S(\Gamma) & \xrightarrow[\pi_{-i}]{} & S(\Gamma_{-i}). \end{array}$$

The map ρ_1 splits the sequence. For exactness, consider the composite

$$H_1(B_2(\Gamma_1)) \oplus H_1(B_2(\Gamma_{-1})) \to H_1(B_2(\Gamma)) \to H_1(B_2(\Gamma_1)) \oplus H_1(B_2(\Gamma_{-1})),$$

where the second map is $(\pi_1, -\pi_{-1})$. Any map in the kernel of the first map is a fortiori in the kernel of the composition. Now writing ι_* for any map on homology induced by an inclusion, and using the commuting diagram above, this composition is given by

$$(\beta_1, \beta_{-1}) \mapsto (\beta_1 - \iota_*(\rho_{-1})_*(\beta_{-1}), \beta_{-1} - \iota_*(\rho_1)_*(\beta_1)).$$

This expression vanishes if and only if $\beta_1 = \iota_*(\rho_1)_*(\beta_1)$, in which case (β_1, β_{-1}) is the image of $(\rho_1)_*(\beta_1)$, showing exactness. □

References

[Abr00] A. Abrams, *Configuration spaces and braid groups of graphs*, Ph.D. thesis, UC Berkeley, 2000.

[ADCK] B. H. An, G. C. Drummond-Cole, and B. Knudsen, *Edge stabilization in the homology of graph braid groups*, 2018, arXiv:1806.05585.

[AF15] D. Ayala and J. Francis, *Factorization homology of topological manifolds*, J. Topology 8 (2015), no. 4, 1045–1084.

[AFT17] D. Ayala, J. Francis, and H. Tanaka, *Factorization homology of stratified spaces*, Selecta Math. (N.S.) 23 (2017), no. 1, 293–362.

[AP17] B. H. An and H. W. Park, *On the structure of braid groups on complexes*, Topology Appl. 226 (2017), no. 1, 86–119.

[Arn69] V. Arnold, *The cohomology ring of the group of dyed braids*, Mat. Zametki 5 (1969), 227–231.

[BC88] C.-F. Bödigheimer and F. Cohen, *Rational cohomology of configuration spaces of surfaces*, Algebraic Topology and Transformation Groups, Lecture Notes in Math., vol. 1361, Springer, 1988, pp. 7–13.

[BCT89] C.-F. Bödigheimer, F. Cohen, and L. Taylor, *On the homology of configuration spaces*, Topology 28 (1989), 111–123.

[Bel04] P. Bellingeri, *On presentations of surface braid groups*, J. Algebra 274 (2004), no. 2, 543–563.

[BM08] A. Bondy and M. R. Murty, *Graph theory*, Grad. Texts in Math., vol. 244, Springer-Verlag, 2008.

[Böd87] C.-F. Bödigheimer, *Stable splittings of mapping spaces*, Algebraic topology (H. Miller and D. Ravenel, eds.), Lecture Notes in Math., vol. 1286, Springer, 1987, pp. 174–187.

[CE80] W. H. Cunningham and J. Edmonds, *A combinatorial decomposition theory*, Canad. J. Math. 32 (1980), 734–765.

[CLM76] F. Cohen, T. Lada, and J. P. May, *The homology of iterated loop spaces*, Lecture Notes in Math., no. 533, Springer, 1976.

[DI04] D. Dugger and D. Isaksen, *Topological hypercovers and $\mathbb{A}^1$-realizations*, Math. Z. 246 (2004), no. 4, 667–689.

[Dug17] D. Dugger, *A primer on homotopy colimits*, preprint, 2008–2017: http://pages.uoregon.edu/ddugger/hocolim.pdf.

[Far03] M. Farber, *Topological complexity of motion planning*, Discrete Comput. Geom. 29 (2003), 211–221.

[Far05] M. Farber, *Collision free motion planning on graphs*, Algorithmic Foundations of Robotics VI (M. Erdmann, M. Overmars, D. Hsu, and F. van der Stappen, eds.), Springer Tracts Adv. Robot., vol. 17, Springer, 2005, pp. 123–138.

[Far06] D. Farley, *Homology of tree braid groups*, preprint, 2006: http://www.users.miamioh.edu/farleyds/grghom.pdf.

[FN62] E. Fadell and L. Neuwirth, *Configuration spaces*, Math. Scand. 10 (1962), 111–118.

[For98] R. Forman, *Morse theory for cell complexes*, Adv. Math. 134 (1998), no. 1, 90–145.

[FS05] D. Farley and L. Sabalka, *Discrete Morse theory and graph braid groups*, Alg. Geom. Topol. 5 (2005), 1075–1109.

[FS08] D. Farley and L. Sabalka, *On the cohomology rings of tree braid groups*, J. Pure Appl. Algebra 212 (2008), no. 1, 53–71.

[FS12] D. Farley and L. Sabalka, *Presentations of graph braid groups*, Forum Math. 24 (2012), 827–860.

[FT00] Y. Félix and J.-C. Thomas, *Rational Betti numbers of configuration spaces*, Topology Appl. 102 (2000), 139–149.

[Gal01] Ś. Gal, *Euler characteristic of the configuration space of a complex*, Colloq. Math. 89 (2001), no. 1, 61–67.

[Ghr02] R. Ghrist, *Configuration spaces and braid groups on graphs in robotics*, Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998), AMS/IP Stud. Adv. Math., vol. 24, Amer. Math. Soc., 2002, pp. 29–40.

[Gin13] G. Ginot, *Notes on factorization algebras, factorization homology and applications*, 2013, preprint: https://arxiv.org/abs/1307.5213.

[GJ09] P. Goerss and J. F. Jardine, *Simplicial homotopy theory*, Modern Birkhäuser Classics, Birkhäuser, 2009.

[GK98] R. Ghrist and D. Koditschek, *Safe cooperative robot patterns via dynamics on graphs*, Robotics Research: The Eighth International Symposium (Y. Shirai and S. Hirose, eds.), Springer, 1998, pp. 81–92.

[Hir02] P. Hirschhorn, *Model categories and their localizations*, Math. Surveys Monogr., vol. 99, Amer. Math. Soc., Providence, Rhode Island, 2002.

[HKRS14] J. M. Harrison, J. P. Keating, J. M. Robbins, and A. Sawicki, *n-particle quantum statistics on graphs*, Comm. Math. Phys. 330 (2014), 1293–1326.

[IMW16] S. Ivanov, R. Mikhailov, and J. Wu, *On nontriviality of certain homotopy groups of spheres*, Homology Homotopy Appl. 18 (2016), no. 2, 337–344.

[KKP12] J. H. Kim, H. K Ko, and H. W. Park, *Graph braid groups and right-angled Artin groups*, Trans. Amer. Math. Soc. 364 (2012), 309–360.

[Knu17] B. Knudsen, *Betti numbers and stability for configuration spaces via factorization homology*, Alg. Geom. Topol. 17 (2017), no. 5, 3137–3187.

[Kon92] M. Kontsevich, *Intersection theory on the moduli space of curves and matrix Airy function*, Comm. Math. Phys. 147 (1992), no. 1, 1–23.

[KP12] K. H. Ko and H. W. Park, *Characteristics of graph braid groups*, Discrete Comput. Geom. 48 (2012), no. 4, 915–963.

[KS06] M. Kashiwara and P. Schapira, *Categories and sheaves*, Grundlehren Math. Wiss., vol. 332, Springer, 2006.

[LS05] R. Longoni and P. Salvatore, *Configuration spaces are not homotopy invariant*, Topology 44 (2005), no. 2, 375–380.

[Lur09] J. Lurie, *Higher topos theory*, Ann. of Math. Stud., Princeton University Press, 2009.

[Lüt14] Daniel Lütgehetmann, *Configuration spaces of graphs*, Master's thesis, Freie Universität Berlin, 2014.

[LW69] A. Lundell and S. Weingram, *The topology of CW complexes*, The University Series in Higher Mathematics, vol. 4, Springer, 1969.

[McD75] D. McDuff, *Configuration spaces of positive and negative particles*, Topology 14 (1975), 91–107.

[Men27] K. Menger, *Zur allgemeinen Kurventheorie*, Fund. Math. 10 (1927), no. 1, 96–115.

[MS17] T. Maciążek and A. Sawicki, *Homology groups for particles on one-connected graphs*, J. Math. Phys. 58 (2017), no. 6.

[Ram18] E. Ramos, *Stability phenomena in the homology of tree braid groups*, Alg. Geom. Topol. 18 (2018), no. 4, 2305–2337.

[Rie14] E. Riehl, *Categorical homotopy theory*, New Mathematical Monographs, vol. 24, Cambridge University Press, 2014.

[Sab09] L. Sabalka, *On rigidity and the isomorphism problem for tree braid groups*, Groups Geom. Dyn. 3 (2009), no. 3, 469–523.

[Seg73] G. Segal, *Configuration-spaces and iterated loop-spaces*, Invent. Math. 21 (1973), 213–221.

[Smi67] L. Smith, *Homological algebra and the Eilenberg–Moore spectral sequence*, Trans. Amer. Math. Soc. 129 (1967), no. 1, 58–93.

[Świ01] Jacek Świątkowski, *Estimates for homological dimension of configuration spaces of graphs*, Colloq. Math. 89 (2001), no. 1, 69–79.

[Tot96] B. Totaro, *Configuration spaces of algebraic varieties*, Topology 35 (1996), no. 4, 1057–1067.

[Tut66] W. T. Tutte, *Connectivity in graphs*, Mathematical Expositions, vol. 15, Toronto University Press, 1966.

[Wag37] K. Wagner, *Über eine Eigenschaft der ebenen Komplexe*, Math. Ann. 114 (1937), no. 1, 570–590.

[WG] J. D. Wiltshire-Gordon, *Models for configuration space in a simplicial complex*, to appear in Colloq. Math.

Byung Hee An
Center for Geometry and Physics
Institute for Basic Science
Pohang 37673
Republic of Korea
anbyhee@gmail.com

Gabriel C. Drummond-Cole
Center for Geometry and Physics
Institute for Basic Science
Pohang 37673
Republic of Korea
gabriel@ibs.re.kr

Ben Knudsen
Mathematics Department
Northeastern University
Boston 02115
USA
b.knudsen@northeastern.edu