Freie Universität Berlin

Representation Stability for Configuration Spaces of Graphs

Dissertation zur Erlangung des Grades eines Doktors der Naturwissenschaften (Dr. rer. nat.) am Fachbereich für Mathematik und Informatik der Freien Universität Berlin

von

Daniel Milan Lütgehetmann

Berlin
2017

Erstgutachter: Prof. Dr. Holger Reich (Freie Universität Berlin)
Zweitgutachter: Prof. Dev Prakash Sinha (University of Oregon)
Drittgutachter: Prof. Dr. Elmar Vogt (Freie Universität Berlin)

Tag der Disputation: 06. Oktober 2017

Diese Arbeit wurde gefördert von der Deutschen Forschungsgemeinschaft durch die Graduiertenschule *Berlin Mathematical School*.

Acknowledgements

First and foremost I would like to thank my advisor Holger Reich for introducing me to the topic of this thesis and for his immense support and encouragement during my studies. I am thankful to him and Elmar Vogt for countless interesting and fruitful discussions about various topics in topology and beyond.

Special thanks go to Dev Sinha for sharing his thoughts on configuration spaces with me and for introducing me to Safia Chettih. Joint work with Safia led to some of the results in this thesis.

I would also like to thank Rachael Boyd, Safia Chettih, Daniela Egas Santander, Filipp Levikov, Peter Patzt, Nils Prigge and Mark Ullmann for all the mathematical discussions we had over the years.

Lastly, I would like to thank my parents Katharina and Ralph, my brother Julian and my girlfriend Charlotte Tumescheit for their endless support, without which I would not have been able to write this dissertation.

Daniel Lütgehetmann
Berlin, August 2017

Contents

Introduction

For a finite graph G and a natural number n, we are interested in the n-*th ordered configuration space of* G. This is the space of n distinguishable, non-colliding particles in G:

$$\mathrm{Conf}_n(G) := \{(x_1, \ldots, x_n) \mid x_i \neq x_j \text{ for } i \neq j\} \subset G^{\times n}.$$

In this dissertation we investigate the singular homology of these spaces.

We construct a concrete generating set for the *first* homology of $\mathrm{Conf}_n(G)$ and if G is a tree, we see that the homology is generated by products of these 1-classes. For general graphs this is far from true, which we observe by describing the following example:

$\mathrm{Conf}_3\left(\right) \simeq$ 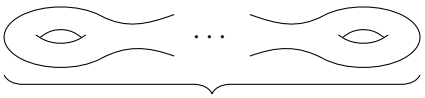

Genus 13

Following the description of these homology groups, we describe how they change if we alter either the number of particles or the base graph G. We consider fixing the number of particles and stabilizing the graph with a graph stabilization process, which is made precise below. We prove that in many cases the sequence of homology groups that arise in this setting has an eventually uniform description (in the sense of representation stability as introduced in [CF13]).

If we instead consider a fixed graph G and vary the number of particles, forgetting the last particle gives us a map

$$\mathrm{Conf}_n(G) \to \mathrm{Conf}_{n-1}(G)$$

for each $n \geqslant 2$. Applying *cohomology* this gives sequences of groups

$$\cdots \to H^i(\mathrm{Conf}_{n-1}(G)) \to H^i(\mathrm{Conf}_n(G)) \to H^i(\mathrm{Conf}_{n+1}(G)) \to \cdots.$$

In [Lü14] we showed that for each G there exists an i such that this sequence is not representation stable. In this dissertation we show that for $i = 1$ and G a 3-vertex connected graph this sequence *does* stabilize.

Throughout this dissertation we will use three main ingredients in order to describe the homology of configuration spaces of graphs:

- configuration spaces with sinks,
- a combinatorial model of these generalized configuration spaces and
- a Mayer-Vietoris spectral sequence comparison argument.

Configuration spaces with sinks — as introduced in [CL16] — are defined for a set of sinks $Z \subset G$ as

$$\mathrm{Conf}_n(G, Z) := \left\{(x_1, \ldots, x_n) \,|\, x_i \neq x_j \text{ or } x_i = x_j \in Z \text{ for } i \neq j\right\} \subset G^{\times n}.$$

For $Z = \varnothing$ this specializes to ordinary configuration spaces. The introduction of sinks allows quotients of the base space to be taken: if $H \subset G$ is a subspace then the quotient map $G \to G/H$ does *not* induce a map on configuration spaces, because any two particles in H would be sent to the same point in G/H. It does, however, induce a map

$$\mathrm{Conf}_n(G) \to \mathrm{Conf}_n(G/H, H/H).$$

This allows us to collapse subgraphs of G and investigate which homology classes survive this process.

The introduction of a combinatorial model for configuration spaces with sinks allows explicit calculations for small graphs, which we later use to describe the homology for bigger graphs.

These generalized configuration spaces also give a geometric decomposition of (parts of) the second page of the Mayer-Vietoris spectral sequence for special choices of open covers of $\mathrm{Conf}_n(G)$. These decompositions allow us to describe parts of the infinity pages, leading to the results described below.

This work has been published in parts as the arXiv preprints [CL16] and [Lü17]. All results of the joint work with Safia Chettih are cited as such.

Statement of the results

Our study of configuration spaces of graphs consists of four parts. We now describe the main contents of these parts.

Chapter 1

We recall the definition of ordered and unordered configuration spaces and their generalization with added sinks. After elaborating on basic properties of these spaces we construct a combinatorial model of the configuration space of n particles in the graph G with sink set Z, denoted by $\mathrm{Conf}_n(G, Z)$, for any finite graph G and Z a subset of the set of vertices of G. This model is a generalization of the model for $\mathrm{Conf}_n(G)$ constructed in [Św01] and [Lü14]. We then use this combinatorial model to compute the homology explicitly for small graphs, for example star graphs. This latter case was already computed by Ghrist in [Ghr01], we recall it due to it being a main ingredient when constructing the generating sets in Chapter 3.

There are other combinatorial models for ordinary configuration spaces, the most well-known of which is constructed by Abrams in [Abr00]. Combined with discrete Morse theory (see [FS05]), the Abrams model can be used to investigate the homology of configuration spaces of graphs. The model in this thesis has an advantage over the Abrams model due to it being much smaller, both in terms of cell numbers and dimension: it has dimension $\min\{n, |V(G)|\}$, whereas the Abrams model has dimension n.

Chapter 2

We give a short proof of the Mayer-Vietoris spectral sequence for an open cover of a space and describe the construction of the boundary maps. Following this we associate to an open cover of a space X an open cover of $\mathrm{Conf}_n(X)$ and describe the corresponding Mayer-Vietoris spectral sequence. For special choices of open covers of a graph, we then describe the entries on the first page in terms of configuration spaces of the open sets in the cover of the base space.

We then introduce the sink comparison argument, which is a technique identifying parts of the E^2-page of such Mayer-Vietoris spectral sequences for *two different graphs*. If we understand the homology for one of these two graphs, then this allows us to transfer that knowledge to the E^2-page for the other graph. To illustrate this technique we compute the homology of configurations in the H-shaped graph, using only the explicit calculations from Chapter 1.

In [MS17] the authors also used Mayer-Vietoris arguments to describe the homology of *unordered* configuration spaces of graphs, which for the first homology leads to an analogous result. Their method of computation is of a different flavor: they perform concrete calculations of Mayer-Vietoris long exact sequences and their boundary maps, whilst we decompose the spectral sequence into parts, which we then compute via geometric considerations.

Chapter 3

A finite graph G is a *tree with loops* if it can be constructed by starting with a tree and taking the iterated wedge with copies of S^1 for different choices of base points. For these graphs we have a much better understanding of the homology of $\mathrm{Conf}_n(G)$ than for general graphs. We prove that this homology is torsion-free:

Theorem A ([CL16, Theorem A, p. 2])**.** *Let G be a tree with loops and let n be a natural number. Then the integral homology $H_q(\mathrm{Conf}_n(G);\mathbb{Z})$ is torsion-free for each $q \geqslant 0$.*

Furthermore, we describe a concrete generating set in terms of *basic classes*: Basic classes are classes in the first homology of configuration spaces of either star graphs, the circle S^1 or the H-graph, considered as classes in $\mathrm{Conf}_n(G)$ by embedding those graphs into G (see Definition 3.1).

Theorem B ([CL16, Theorem B, p. 2])**.** *Let G be a tree with loops and let n be a natural number. Then the homology of $\mathrm{Conf}_n(G)$ is generated by products of disjoint basic classes.*

We prove these results simultaneously by investigating a particular Mayer-Vietoris spectral sequence arising from gluing a graph with one essential vertex (meaning it has valence at least three) to a tree with loops. By explicitly identifying the homology of configurations in graphs with one essential vertex, we can compute the infinity page of that spectral sequence.

Similar descriptions of the homology were obtained for unordered configurations: in [MS17] the authors describe a generating system for the first homology of unordered configurations consisting of basic classes as defined above and additionally they give concrete formulas for the Euler characteristic in terms of invariants of the graph. Farley showed in [Far06] that the homology groups of unordered configurations in any tree are free and he gives concrete rank calculations for specific graphs. For non-planar graphs, however, Kim, Ko and Park showed in [KKP12] that the unordered configuration spaces have elements of order 2 in their first homology groups. We do not expect this to be the case for ordered configurations, and indeed for the examples of graphs producing torsion in the unordered case in [KKP12] (the complete graphs K_5 and $K_{3,3}$) one can compute that the *ordered* configuration space has torsion-free homology.

For ordered configuration spaces of only two particles, there are results which are similar to the theorems above: in [BF09] and [FH10] the Betti numbers and generators of the homology groups of the spaces $\mathrm{Conf}_2(G)$ are determined for various classes of graphs G. In [Che16] generators for $H_1(\mathrm{Conf}_2(T))$ for finite trees T are given in terms of basic classes as above.

In [Ram17] Ramos considered trees where *all vertices* are sinks. He proved torsion-freeness of their homology and computed the homological dimension of those spaces.

Following this, we describe generators of the *first* homology of configuration spaces of arbitrary graphs:

Theorem C ([CL16, Theorem C, p. 3]). *If* G *is any finite graph and* n *a natural number, then the first homology group* $H_1(\mathrm{Conf}_n(G))$ *is generated by basic classes.*

The proof of this statement uses very similar techniques as the proofs of Theorem A and Theorem B. The naive generalization of Theorem B to general graphs is false, which we illustrate by describing a configuration space which is homotopy equivalent to a surface of genus 13. The fact that the homology is not in general generated by products was already known before: Abrams and Ghrist showed in [AG02] that the configuration space of two particles in the complete graph on 5 vertices is homotopy equivalent to an orientable surface of genus 6. Our example has the advantage that it can be modified to give counter-examples in all dimensions bigger than 2.

In the final part of the chapter we prove the statements above for the case when arbitrary subsets of the vertex set are turned into sinks. This has two purposes: it completes the picture for the more general case of configuration spaces with sinks and the technical tools developed in order to prove the statements are required in Chapter 4. More precisely, we prove the following:

Theorem D ([CL16, Theorem D, p. 3]). *Let* G *be a finite graph and let* Z *be any subset of the vertex set. Then the first homology of* $\mathrm{Conf}_n(G, Z)$ *is generated by basic classes. If* G *is a tree with loops, then* $H_*(\mathrm{Conf}_n(G, Z))$ *is free and generated by products of basic classes.*

Chapter 4

After reviewing the concept of representation stability and **FI**-modules (as defined in [CF13] and [CEF15]) we describe the following construction of graph stabilization. Given three graphs G_0, G_1 and K_1 such that K_1 is a subgraph of G_0 and G_1, we denote by G_k the quotient of the disjoint union of G_0 and k copies of G_1 identifying all copies of K_1 with each other. The symmetric group Σ_k acts on G_k by permuting the copies of G_1, so we can ask the question whether for fixed natural numbers n and i the sequence of representations of symmetric groups

$$\cdots \to H_i(\mathrm{Conf}_n(G_{k-1})) \to H_i(\mathrm{Conf}_n(G_k)) \to H_i(\mathrm{Conf}_n(G_{k+1})) \to \cdots$$

is representation stable.

More generally, for a base graph G_0 and an ℓ-tuple

$$\underline{\Gamma} = \{(K_1 \subset G_1), \ldots, (K_\ell \subset G_\ell)\}$$

such that each K_i is also a subgraph of G_0, we define an $\mathbf{FI}^{\times \ell}$-space $\mathbf{G}_{\underline{\Gamma}}$, which evaluated at $(j_1, \ldots, j_\ell)$ is given as the quotient of the disjoint union of G_0 and j_i

copies of G_i, identifying for each $1 \leqslant i \leqslant \ell$ all copies of K_i with each other. Here, the category $\mathbf{FI}^{\times \ell}$ is the ℓ-fold product category of the category **FI** of finite sets and injections.

This defines an $\mathbf{FI}^{\times \ell}$-module

$$\mathbf{H}^{\underline{\Gamma}}_{q,n} := H_q(\mathrm{Conf}_n(\mathbf{G}_{\underline{\Gamma}}); \mathbb{Z})$$

for each q, n in $\mathbb{N}$, and we can ask whether this is finitely generated.

We will answer this question for different restrictions on q, n and the graphs.

Theorem E. *If each of the graphs G_i for $0 \leqslant i \leqslant \ell$ is a tree with loops, then $\mathbf{H}^{\underline{\Gamma}}_{q,n}$ is finitely generated in degree $(\zeta, \zeta, \ldots, \zeta)$ for each $q, n \in \mathbb{N}$, where $\zeta = \zeta_{n,q} = \min\{2n, n + 3q\}$.*

This implies immediately that for G_0, G_1 trees with loops, the sequence of representations of symmetric groups above is representation stable.

For general graphs, we restrict ourselves to the first homology and recover an analogous statement:

Theorem F. *For any choice of graphs G_i and K_i the* **FI***-module $\mathbf{H}^{\underline{\Gamma}}_{1,n}$ is finitely generated in degree $(n + 3, n + 3, \ldots, n + 3)$ for each $n \in \mathbb{N}$.*

The proofs of both theorems are based on the knowledge of generators for the homology from the previous chapter and explicit constructions for basic classes.

We then prove finite generation for the second and third homology of the **FI**-space $\mathrm{Conf}_n(B_\bullet, Z)$, where B_k is the *banana graph on k edges*, i.e. the graph given by two vertices v and w connected via k edges, and $Z \subset \{v, w\}$.

Theorem G. *For each $n \in \mathbb{N}$ and $q \in \mathbb{N}$ the* **FI***-module $H_q(\mathrm{Conf}_n(B_\bullet))$ is finitely generated in degree $n + 6$. The* **FI***-module $H_q(\mathrm{Conf}_n(B_\bullet, Z))$ is finitely generated in the same degree for any $Z \subset \{v, w\}$ and $q \leqslant 3$.*

The proof of this statement proceeds by reducing it to the case $Z = \{v, w\}$ and then proving that case. As it turns out, the result for $Z = \{v, w\}$ can more generally be used to show finite generation for $\mathbf{H}^{\underline{\Gamma}}_{2,n}$ for arbitrary $\underline{\Gamma}$:

Theorem H. *For any choice of graphs K_i, G_i and $n \in \mathbb{N}$, the $\mathbf{FI}^{\times \ell}$-module $\mathbf{H}^{\underline{\Gamma}}_{2,n}$ is finitely generated in degree $(n + 6, \ldots, n + 6)$.*

In the final part of this chapter we investigate the **FI**-module $H^i(\mathrm{Conf}_\bullet(G); A)$ for an abelian group A and a finite graph G. An injection $T \hookrightarrow S$ determines a map

$$H^i(\mathrm{Conf}_T(G); A) \to H^i(\mathrm{Conf}_S(G); A)$$

induced by precomposition. In [Lü14] we showed that there exists an $i > 0$ such that this **FI**-module is *not* finitely generated. Furthermore, Ramos shows in [Ram16b]

that the rank of the homology groups for *unordered* configuration spaces of *trees* is non-constant in every degree (except zero). This also implies that the ordered configuration space cannot be representation stable in any degree.

The reason why representation stability fails for trees is that for a given class, there are many options for changing the positions of the non-moving particles, and in trees many of these give distinct homology classes. If our graph has enough distinct paths connecting different vertices, however, all the different positions of the non-moving particles will give the same homology class. A graph is k-*vertex connected* if for every pair of vertices v, w there are k paths connecting v and w which are disjoint (except for the two ends of the paths). For such graphs, we can prove representation stability for $i = 1$:

Theorem I. *Let* G *be a finite 3-vertex connected graph with at least four essential vertices and without self-loops. Let* A *be an abelian group such that* $H_1(\mathrm{Conf}_2(G); A)$ *is torsion-free. Then* $H_1(\mathrm{Conf}_n(G); A)$ *is torsion-free for all* n *and the* **FI***-module* $H^1(\mathrm{Conf}_\bullet(G); A)$ *is finitely generated in degree 2. In particular, the sequence* $n \mapsto H^1(\mathrm{Conf}_n(G); \mathbb{Q})$ *induced by forgetting the last particle is representation stable and its dimension is eventually polynomial in* n.

For 3-vertex connected graphs, Theorem I recovers the result of Ko and Park in [KP12], which shows that for *unordered configurations* the sequence $n \mapsto H_1(\mathrm{UConf}_n(G))$ satisfies homological stability. In fact, they prove this for all 2-vertex connected graphs and we expect the analogous result for ordered configuration spaces also to be true in this bigger generality. This alongside the case $i > 1$ will be the subject of further studies.

Stability for configuration spaces of graphs with respect to the number of particles was also investigated by Ramos: in [Ram16b] the author shows that the ranks of the homology groups of unordered configurations in a tree are polynomial in the number of particles. In [Ram17] he proves that the homology groups of configuration spaces in graphs where *all vertices* are sinks satisfy "generalized representation stability", a term introduced in [Ram16a].

Chapter 1

Configuration spaces of graphs and their combinatorial model

In this chapter we introduce configuration spaces of graphs and discuss their combinatorial models. For the computations in later chapters, we will need a modified version of configuration spaces, namely configuration spaces with sinks. We define this generalization and extend the combinatorial model to this broader class of spaces.

1.1 Configuration spaces

Before we focus on graphs, let us define configuration spaces in general:

Definition 1.1. Let X be a topological space and let S be a finite set. Define the *configuration space of the particles* S *in* X as the space of injective maps $S \hookrightarrow X$, i.e.

$$\mathrm{Conf}_S(X) := \{\iota \colon S \hookrightarrow X\} \subset \mathrm{map}(S, X)$$

with the subspace topology of $\mathrm{map}(S, X)$. For $n \in \mathbb{N}$ we write $\underline{\mathbf{n}} := \{1, \dots, n\}$ and call

$$\mathrm{Conf}_n(X) := \mathrm{Conf}_{\underline{\mathbf{n}}}(X)$$

the n*-th ordered configuration space of* X. △

This construction is functorial in S and X if we restrict ourselves to injective maps: an injective map $S \hookrightarrow T$ induces a map $\mathrm{Conf}_T(X) \to \mathrm{Conf}_S(X)$ by precomposition; a continuous injection $X \hookrightarrow Y$ determines an injection $\mathrm{Conf}_S(X) \hookrightarrow \mathrm{Conf}_S(Y)$ by postcomposition. In particular, an isomorphism $S \xrightarrow{\cong} T$ of finite sets induces a homeomorphism $\mathrm{Conf}_T(X) \xrightarrow{\cong} \mathrm{Conf}_S(X)$. Denote for a finite set S by Σ_S the symmetric group on S:

$$\Sigma_S := \{S \hookrightarrow S\} \subset \mathrm{map}(S, S).$$

The right action of the symmetric group Σ_S on $\mathrm{Conf}_S(X)$ by precomposition is free. We again write $\Sigma_n := \Sigma_{\underline{n}}$ and thus get an action of Σ_n on $\mathrm{Conf}_n(X)$.

Definition 1.2. Let X be a topological space and let $n \in \mathbb{N}$, then we define the n*-th unordered configuration space of* X as

$$\mathrm{UConf}_n(X) := \mathrm{Conf}_n(X)/\Sigma_n. \qquad \triangle$$

Remark 1.3. For a disjoint union of spaces X and Y the configuration space of a finite set of particles S in $X \sqcup Y$ can be described as follows:

$$\mathrm{Conf}_S(X \sqcup Y) \cong \coprod_{S_X \sqcup S_Y = S} \mathrm{Conf}_{S_X}(X) \times \mathrm{Conf}_{S_Y}(Y).$$

Therefore, we will mostly restrict our investigations to configuration spaces of connected spaces. $\triangle$

One important property of configuration spaces is that a map $f\colon X \to Y$ induces a map on configuration spaces *only if* f *is injective* because otherwise, f could map two particles of a specific configuration to the same point in Y. In particular, a quotient map $X \to X/\sim$ does not induce a map of configuration spaces. In order to define such a map we introduced in [CL16] the notion of configuration spaces *with sinks*.

Definition 1.4. Let X be a topological space, let S be a finite set and let $Z \subset X$ be a subspace. Then the *ordered configuration space of* S *in* X *with sink set* Z is defined as

$$\mathrm{Conf}_S(X, Z) := \left\{\iota\colon S \to X \,\middle|\, \text{for all } s \neq s' \text{ either } \iota(s) \neq \iota(s') \text{ or } \iota(s) = \iota(s') \in Z\right\},$$

endowed with the subspace topology of $\mathrm{map}(S, X)$. Notice that we have $\mathrm{Conf}_S(X, \varnothing) = \mathrm{Conf}_S(X)$. $\triangle$

The set of sinks defines special parts of the space X in which particles are allowed to collide. If we now have an arbitrary map of pairs $f\colon (X, Z_X) \to (Y, Z_Y)$, i.e. $f(Z_X) \subset Z_Y$, and Z_Y contains the set of points with at least two preimages under f, then f induces a map

$$\mathrm{Conf}_S(f)\colon \mathrm{Conf}_S(X, Z_X) \to \mathrm{Conf}_S(Y, Z_Y).$$

Such maps can be used to detect homology classes by collapsing parts of X and turning those parts into sinks. Notice that we still have a right action of Σ_S on $\mathrm{Conf}_S(X, Z)$, but this action is not free anymore if Z is non-empty.

Configuration spaces with sinks are functorial under the inclusion of sink sets, i.e. for $Z \subset Z'$ we get an injective map

$$\mathrm{Conf}_S(X, Z) \hookrightarrow \mathrm{Conf}_S(X, Z').$$

For $Z = X$ we get $\mathrm{Conf}_S(X, X) = \mathrm{map}(S, X)$, so the set of sinks interpolates between the ordinary configuration space of X and the space of arbitrary maps $S \to X$:

$$\mathrm{Conf}_S(X) = \mathrm{Conf}_S(X, \varnothing) \hookrightarrow \mathrm{Conf}_S(X, Z) \hookrightarrow \mathrm{Conf}_S(X, X) = \mathrm{map}(S, X)$$

Definition 1.5. A *space with sinks* is a tuple of a topological space X and a subspace Z. These spaces form a category $\mathbf{Top}^{\mathbf{sink}}$, which we call the *category of spaces with sinks*, where morphisms between (X, Z_X) and (Y, Z_Y) are given by continuous maps $f\colon X \to Y$ such that Z_Y contains $f(Z_X)$ and all points of Y having strictly more than one preimage under f. This category has an initial object $(\varnothing, \varnothing)$ and a terminal object $(\mathrm{pt}, \mathrm{pt})$. Each morphism set $\mathrm{map}_{\mathbf{Top}^{\mathbf{sink}}}((X, Z_X), (Y, Z_Y))$ is endowed with the subspace topology of $\mathrm{map}(X, Y)$. △

This definition is tailored in such a way that Conf_S can be seen as a (continuous) functor

$$\begin{aligned} \mathrm{Conf}_S\colon \mathbf{Top}^{\mathbf{sink}} &\to \mathbf{Top} \\ (X, Z) &\mapsto \mathrm{Conf}_S(X, Z). \end{aligned}$$

If **FI** is the category of finite sets and injections (see Section 4.1), then we get the following functor:

$$\mathrm{Conf}_\bullet\colon \mathbf{FI} \times \mathbf{Top}^{\mathbf{sink}} \to \mathbf{Top}$$

Definition 1.6. A *homotopy of maps between spaces with sinks* (X, Z_X) *and* (Y, Z_Y) is a continuous map

$$f\colon [0, 1] \to \mathrm{map}_{\mathbf{Top}^{\mathbf{sink}}}((X, Z_X), (Y, Z_Y)).$$ △

Proposition 1.7. *A homotopy of maps between spaces with sinks induces a homotopy of the corresponding maps between configuration spaces with sinks.*

Proof. This follows straight from the definition of homotopy by postcomposing with Conf_S, the homotopy is given by

$$[0, 1] \to \mathrm{map}_{\mathbf{Top}^{\mathbf{sink}}}((X, Z_X), (Y, Z_Y)) \to \mathrm{map}_{\mathbf{Top}}(\mathrm{Conf}_S(X, Z_X), \mathrm{Conf}_S(Y, Z_Y)).$$ □

Corollary 1.8. *If* $f \in \mathrm{map}_{\mathbf{Top}^{\mathbf{sink}}}((X, Z_X), (Y, Z_Y))$ *is a homotopy equivalence of spaces with sinks, then the induced map* $\mathrm{Conf}_S(f)\colon \mathrm{Conf}_S(X, Z_X) \to \mathrm{Conf}_S(Y, Z_Y)$ *is a homotopy equivalence of topological spaces.*

Definition 1.9. Let X be a topological space and let $n \in \mathbb{N}$. The n*-th unordered configuration space with sinks of* (X, Z) is defined as

$$\mathrm{UConf}_n(X, Z) := \mathrm{Conf}_n(X, Z)/\Sigma_n.$$ △

In this thesis we will almost exclusively be interested in *ordered* configuration spaces, so for the sake of brevity, we omit the word "ordered" whenever there is no possibility of confusion.

Definition 1.10. For a space with sinks (X, Z) and finite sets $S' \subset S$ we define the projection to the particles S' by

$$\begin{aligned} \pi_{S'} \colon \operatorname{Conf}_S(X, Z) &\to \operatorname{Conf}_{S'}(X, Z) \\ (\iota \colon S \hookrightarrow X) &\mapsto (\iota \circ \operatorname{inc} \colon S' \hookrightarrow S \hookrightarrow X), \end{aligned}$$

where $\operatorname{inc} \colon S' \hookrightarrow S$ is the inclusion. For $S' = \{s\}$ we also write $\pi_s = \pi_{\{s\}}$. △

1.2 Configurations in graphs

Let us now define the ambient spaces in whose configuration spaces we are interested in, namely graphs and graphs with sinks.

Definition 1.11. A *graph* G is a topological space with the structure of a 1-dimensional CW complex. We call the zero-dimensional cells *vertices* and denote the set of vertices by $V(G)$. One-dimensional cells are called *edges* and the set of edges is denoted by $E(G)$. All graphs considered in this thesis are assumed to be finite, meaning that the sets $V(G)$ and $E(G)$ are finite. A vertex is called *essential* if it has valence at least three.

A *graph with sinks* (G, Z) is a tuple consisting of a graph G and a subspace $Z \subset G$. In this thesis we will only be interested in the case where Z is a union of closed edges and vertices of G, so we will always assume that. The category of finite graphs with sinks $\mathbf{Graph}^{\mathbf{sink}}$ is the full subcategory of $\mathbf{Top}^{\mathbf{sink}}$ consisting of all such graphs with sinks, i.e. all those objects (X, Z) where X has the structure of a 1-dimensional CW complex and Z is a union of closed edges and vertices. We will see below that the most interesting case is when Z only consists of vertices, so if not stated otherwise, we will assume that Z does not contain any edges. △

We will always assume that we have the smallest combinatorial model of each graph G, namely the one without vertices of valence two (except for the circle S^1, which consists of precisely one edge and one vertex of valence two).

As it turns out, there is a special class of graphs whose configuration space homology we can describe rather explicitly, namely the class of trees with loops:

Definition 1.12. Let X, Y and Z be unpointed spaces. We say that Z is *a wedge of* X *and* Y if there exist base points $x \in X$ and $y \in Y$ such that the pointed wedge sum $X \vee Y$ is homeomorphic (after forgetting the base point) to Z. Given $k \geqslant 2$ and unpointed spaces $X_1, \dots, X_k$ we say that Z is *an iterated wedge of* $X_1, \dots, X_k$ if Z is a wedge of X_1 and an iterated wedge of $X_2, \dots, X_k$, where an iterated wedge of two unpointed

spaces is the ordinary wedge from above. There can be multiple spaces Z which are iterated wedges of fixed spaces $X_1, \dots, X_k$, for example every finite tree on k edges is an iterated wedge of k intervals. △

Definition 1.13. A finite connected graph G is called a *tree with loops* if it can be constructed as an iterated wedge of star graphs and copies of S^1. In particular, every finite tree is a tree with loops. △

We will see that the difficulties in understanding the homology of the configuration spaces of a graph G do not depend on the rank of the graph but rather on the vertex connectivity of G. Trees with loops are precisely those graphs which disconnect after removing any single vertex.

As a first simplification, we show that we can collapse contractible edges in the sink set to point shaped sinks.

Proposition 1.14. *Let (G, Z) be a graph with sinks, let $e \subset Z$ be an edge which does not form a self-loop and let S be a finite set. Then the collapse map $\pi\colon (G, Z) \to (G/e, Z/e)$ induces a homotopy equivalence*

$$\mathrm{Conf}_S(\pi)\colon \mathrm{Conf}_S(G, Z) \xrightarrow{\simeq} \mathrm{Conf}_S(G/e, Z/e).$$

Proof. By Corollary 1.8, we only have to show that $(G, Z) \to (G/e, Z/e)$ is a homotopy equivalence of spaces with sinks. A homotopy inverse of $G \to G/e$ is given by collapsing the edges of a small star around the vertex e/e to two segments of an interval and mapping that interval to e, with e/e mapping to the midpoint of e, see Figure 1.1. This gives a well-defined map

$$\phi\colon (G/e, Z/e) \to (G, Z)$$

of spaces with sinks. The compositions $\pi \circ \phi$ and $\phi \circ \pi$ are each homotopic to the identity, which can be seen by starting with the identity and pulling the particles towards the sink edge e or the sink e/e. □

Collapsing such edges eventually yields a graph with sinks (G', Z'), where Z' consists only of vertices and edges forming self-loops. We will see later that sink edges forming self-loops are not that interesting (see Proposition 3.12), so unless stated otherwise, we will always assume that Z consists only of vertices.

For the description of the homology, it will be useful to have the following notion of a *product of disjoint classes*:

Definition 1.15. A homology class $\sigma \in H_q(\mathrm{Conf}_n(G))$ is called the *product of classes* $\sigma_1 \in H_{q_1}(\mathrm{Conf}_{T_1}(G_1))$ *and* $\sigma_2 \in H_{q_2}(\mathrm{Conf}_{T_2}(G_2))$ if it is the image of $\sigma_1 \otimes \sigma_2$ under the map

$$H_q(\mathrm{Conf}_n(G_1 \sqcup G_2)) \to H_q(\mathrm{Conf}_n(G))$$

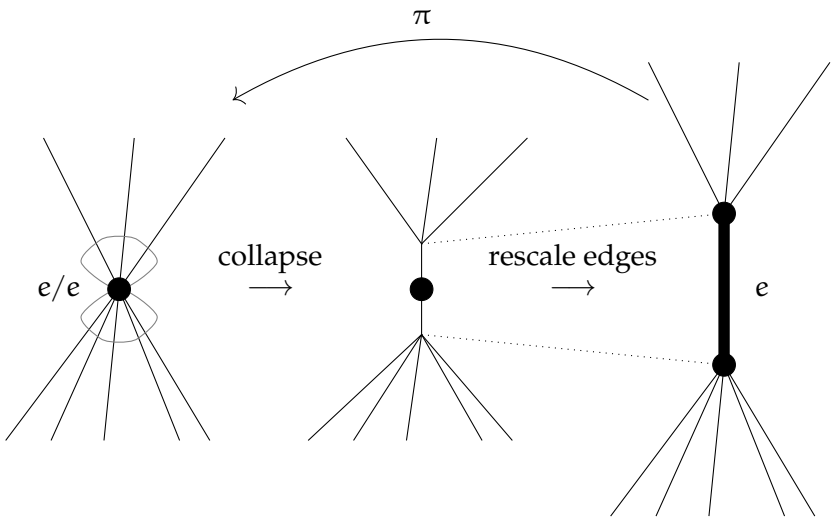

Figure 1.1: The composition of these two maps is the homotopy inverse of the collapse map π.

induced by an embedding $G_1 \sqcup G_2 \hookrightarrow G$.

Analogously, we define iterated products. △

Notice that q_1 or q_2 could be zero, so a product in the q-th homology can be a product with more than q factors.

It will also be useful for subsequent proofs to have a notion for pushing in new particles from the boundary of the graph.

Definition 1.16. Let G be a graph and let e be a leaf. For a finite set S and an element $s \in S$, define the map

$$\iota_{e,s} \colon \mathrm{Conf}_{S-\{s\}}(G, W) \hookrightarrow \mathrm{Conf}_S(G, W)$$

by slightly pushing in the particles on e and putting s onto the univalent vertex of e. △

Notice that the composition $\pi_{S-\{s\}} \circ \iota_{e,s}$ is homotopic to the identity.

1.3 Combinatorial models

For unordered configurations in graphs, Świątkowski introduced a combinatorial model in [Św01]. In [Lü14] we adapted the argument to give a combinatorial model for *ordered* configuration spaces:

Theorem 1.17 ([Lü14, Theorem 2.3, p. iii])**.** *Let* G *be a finite graph and let* $n \in \mathbb{N}$*. Then there exists a finite cube complex which is a deformation retract of* $\mathrm{Conf}_n(G)$*. Its dimension is given by* $\min\{n, |V(G)|\}$*.*

Let us briefly recall what we mean by cube complex.

Definition 1.18 (Cube Complex, [BH99, Definition I.7.32])**.** A cube complex K is the quotient of a disjoint union of cubes $X = \bigsqcup_{\lambda \in \Lambda} [0,1]^{k_\lambda}$ by an equivalence relation $\sim$ such that the quotient map $p\colon X \to X/\sim = K$ maps each cube injectively into K and we only identify faces of the same dimensions by an isometric homeomorphism. △

Remark 1.19. In the original definition by Bridson and Häfliger two cubes cannot be identified along more than one face, so in particular between two vertices there cannot be two distinct 1-cubes connecting them. This, however, happens in the complex we want to describe, so we need this slight generalization. △

In [CL16] we generalized this construction to configuration spaces with sinks. More precisely we defined a deformation retraction $r\colon \mathrm{Conf}_n(G,Z) \to \mathrm{Conf}_n(G,Z)$ such that the image of r has the structure of a finite cube complex. Each axis of such a cube corresponds to the combinatorial movement of one particle. Such a combinatorial movement is either given by the movement from an essential non-sink vertex onto an edge or along a single edge from one sink to the other. Each vertex and each such edge can only be involved in one of those combinatorial movements at the same time, so the dimension of this cube complex will be restricted by the number of essential non-sink vertices and the edges connecting two sinks.

Proposition 1.20 ([CL16, Proposition 2.1, p. 4])**.** *Let* G *be a finite graph, let* $Z \subset V(G)$ *and let* $n \in \mathbb{N}$*. Then there exists a finite cube complex which is a deformation retract of* $\mathrm{Conf}_n(G,Z)$*. Its dimension is given by* $\min\{n, |V(G) - Z| + |E_{\mathrm{sink}}(G)|\}$*, where* $E_{\mathrm{sink}}(G)$ *is the set of edges whose initial* and *terminal vertices are sinks.*

Proof. This proof is the same as the one in [CL16] with some additional details for the description of the cubes. Give G the path metric such that every edge has length 1 and choose arbitrary orientations for all edges. The general idea is now the following: the retraction r only changes the position of particles *inside* (closed) edges of the graph. We move as many particles of a given configuration $\mathbf{x} = (x_1, \ldots, x_n)$ as possible into the sinks, so that $r(\mathbf{x})$ has at most one particle in the interior of any edge incident to a sink. Furthermore, the particles of $r(\mathbf{x})$ on each single edge will be equidistant, except for the outermost particles, which may be closer to the vertices, see Figure 1.2. The main difficulty will be to define for each configuration $\mathbf{x}$ and each edge the parameters $t_e^\iota, t_e^\tau \in [0,1]$ determining the distance from the vertices. Decreasing t_e^ι to zero represents moving the first particle on the edge towards the initial vertex $\iota(e)$ of e. To avoid multiple particles approaching the same vertex, we,

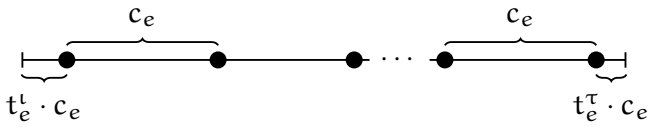

Figure 1.2: Equidistant particles on e.

therefore, require that for any pair of edges $e \neq e'$ with the same initial vertex only one of the two values t_e^ι and $t_{e'}^\iota$ can be strictly smaller than 1.

For fixed $(x_1, \dots, x_n) \in \mathrm{Conf}_n(G, Z)$ we now define the image $r(x)$. Let e be an edge of G and remove the set of particles sitting in the interior of e. This cuts e into intervals, and we denote by ℓ_e^ι and ℓ_e^τ the length of the first and last segment of e, respectively. Let k_e be the number of particles in the interior of e and define δ_e^τ as follows:

$$\delta_e^\tau := \begin{cases} \ell_e^\tau & \text{if } \iota(e) \text{ is a sink} \\ \dfrac{\ell_e^\tau}{\ell_e^\iota} & \text{if } k_e \leqslant 1 \text{ and } \iota(e) \text{ is not a sink} \\ \ell_e^\tau \dfrac{k_e - 1}{1 - \ell_e^\tau - \ell_e^\iota} & \text{else.} \end{cases}$$

This gives a "normalized" distance of the last particle on e from the terminal vertex. Define $\delta_e^\iota := \delta_{-e}^\tau$, where $-e$ is the edge e with opposite orientation.

For each edge e, if

- $\tau(e)$ has valence at least three,
- $\tau(e)$ is not occupied by any particle,
- $\delta_e^\tau < 1$ and
- $\delta_e^\tau < \delta_{e'}^\tau$, for all oriented edges e' with $\tau(e) = \tau(e')$,

define $t_e^\tau := \delta_e^\tau$. Otherwise, define $t_e^\tau := 1$. Again, we set $t_e^\iota := t_{-e}^\tau$.

Given these parameters t_e^τ and t_e^ι for all edges e we now construct the configuration $r((x_1, \dots, x_n))$. The particles on the vertices are not moved by the retraction, so it remains to describe the change of position for the particles in the interior of an edge e. We will not change the order of the particles but only their position within the edge, and to make the description more concise we choose once and for all an isometric identification of each edge e with $[0, 1]$.

If e is not incident to a sink vertex the new position of the j-th vertex on e will be given by $(t_e^\iota + j - 1) \cdot c_e$, where $k_e \geqslant 1$ is the number of particles in the interior of e and $c_e := (t_e^\iota + k_e - 1 + t_e^\tau)^{-1}$ will be the distance between the particles on that

edge. This gives all particles on the edge the same distance and only modifies the distances from the vertices, see Figure 1.2.

If precisely one of the initial and terminal vertices of e is a sink then we can assume that this sink vertex corresponds to $0 \in [0,1]$. All particles on e except the last one are then moved to 0, the last particle is moved to $1 - t_e^\tau \in [0,1]$.

If both the initial and terminal vertex of e are sinks we slide all particles away from $1/2 \in [0,1]$ with speed given by their distance from $1/2$ until at most one particle is left in the interior $(0,1)$ of the interval. This gives a configuration having at most one particle on e and the rest on the sinks.

More precisely, we apply the flow $\Phi\colon [0,1] \times [0,\infty) \to [0,1]$ given by

$$(x,t) \mapsto \frac{1}{2} + \left(x - \frac{1}{2}\right)e^t$$

followed by the collapse $\mathbb{R} \to [0,1]$ mapping $x > 1$ to 1 and $x < 0$ to 0. There are at least two particles in the interior of the edge, so let $z \in [0,1]$ be the position of the particle which is the second closest to $1/2$. We flow along Φ until at most one particle is in $(0,1)$, which will occur when $t = -\ln(|1-2z|)$.

It is straightforward to check that the described map is continuous and a retraction, i.e. satisfies $r^2 = r$. In the description above we only changed the positions of particles on individual edges, so there is an obvious homotopy from the identity to r by just adjusting the positions of the particles on each edge individually.

The image of r has the structure of a cube complex: the 0-cells are configurations where all particles in the interior of each interval cut the interval into pieces of equal length, and additionally no particle is in the interior of an edge with one or two sink vertices. A k-cube is given by choosing such a 0-cell, k distinct particles which are on a vertex and an adjacent edge to each of those k vertices. The k-cube then moves each particle from its original position on the vertex onto the chosen edge (if the other vertex of the edge is not a sink) or onto the sink on the other end of the edge. Such a choice of k movements determines a k-cube if and only if we can realize the movements independently, namely if no two particles move towards the same non-sink vertex and no two particles move along the same edge incident to two sink vertices.

Each direction of the cube corresponds to the movement of one of the particles. The k coordinates of the k-cube $[0,1]^k$ are used as values t_e^τ for the chosen oriented edges e: if a coordinate is 0, then the corresponding particle sits on the vertex. Increasing the coordinate, it moves along the chosen edge e, and at coordinate 1 it arrives at its terminal position, which is either in the interior of e or on the sink $\iota(e)$ (depending on the type of $\iota(e)$). Changing multiple coordinates at the same time moves the different corresponding particles independently. See Figure 1.3 for an

example of a 2-cube.

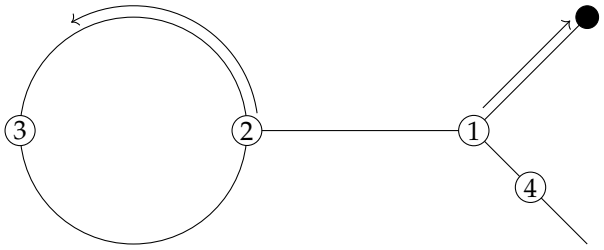

Figure 1.3: A 2-cube in the combinatorial model. Each of the two coordinates of $[0,1]^2$ corresponds to the movement of one of the two particles along the arrow. At the top right corner $(1,1)$, particle 1 will be on the top right sink, and particles 2 and 3 will both be on the circle.

By the restrictions listed above, each tuple $(t_1,\dots,t_k) \in [0,1]^k$ in such a cube determines a point in the configuration space, and each point of the image of r can by definition of the retraction be written uniquely in such a way. In this way the image has the structure of a cube complex.

By the description of the choices involved for finding k-cubes we immediately get the restriction on the dimension. For more details about the general construction of the cube complex (without sinks), see [Lü14]. □

Notice that this already gives an upper bound on the homological dimension of $\mathrm{Conf}_n(G,Z)$: all homology groups in dimension $k > \min\{n, |V(G)-Z| + |E_{\mathrm{sink}}(G)|\}$ are zero. We will later see that if the number of particles is big enough, then this upper bound is sharp. To show this, we first compute the configuration spaces of a few small graphs, which will be the content of the next section.

Remark 1.21. In [Abr00] Abrams constructed a different combinatorial model for configuration spaces of graphs. The idea in his work is to subdivide all edges into $n+1$ smaller edges, where n is the number of particles. He then allows combinatorial movements of individual particles from one vertex to the next one if it is not occupied by another particle.

This model is much bigger than the one described above, both in terms of the number of cells (traversing a single edge in Abrams' model takes a single particle $n+1$ steps, in our model it only takes two steps) and in terms of the dimension (the dimension of Abrams' model is always precisely n, whereas the dimension of the model above stays constant for $n \gg 0$). △

The combinatorial model gives an easy proof of the following fact, which will be useful in later proofs.

Proposition 1.22. *Let* (G, Z) *be a finite graph with sinks and let* e *be a leaf of* G *such that the valence one vertex of* e *denoted by* v *is not a sink. For* $G' = G - \{v\}$ *and a finite set* S *the inclusion*

$$\mathrm{Conf}_S(G', Z) \hookrightarrow \mathrm{Conf}_S(G, Z),$$

induces a homotopy equivalence. Therefore, configuration spaces do not see the difference between open and a closed leaves.

Proof. The deformation retraction from the proof of Proposition 1.20 restricts to the subspace $\mathrm{Conf}_S(G', Z) \subset \mathrm{Conf}_S(G, Z)$ because particles are only pulled away from v as this univalent vertex is not a sink. Since both spaces are homotopy equivalent to the combinatorial model via this deformation, the inclusion is also a homotopy equivalence. □

1.4 Explicit calculations

For small graphs, the explicit combinatorial model from the previous section is sufficient to determine the homology completely.

Proposition 1.23 ([CL16, Proposition 2.5, p.7])**.**

$$\mathrm{H}_i(\mathrm{Conf}_n(I, \varnothing)) = \begin{cases} \mathbb{Z}\Sigma_n & i = 0 \\ 0 & \textit{else} \end{cases}$$

$$\mathrm{H}_i\big(\mathrm{Conf}_n(S^1, \varnothing)\big) = \begin{cases} \mathbb{Z}(\Sigma_n/\mathrm{shift}) \cong \mathbb{Z}^{(n-1)!} & i = 0, 1 \\ 0 & \textit{else} \end{cases}$$

$$\mathrm{H}_i(\mathrm{Conf}_n(I, \{0\})) = \begin{cases} \mathbb{Z} & i = 0 \\ 0 & \textit{else} \end{cases}$$

$$\mathrm{H}_i(\mathrm{Conf}_n(I, \{0, 1\})) = \begin{cases} \mathbb{Z} & i = 0 \\ \mathbb{Z}^{(n-2)2^{n-1}+1} & i = 1 \\ 0 & \textit{else} \end{cases}$$

$$\mathrm{H}_i\big(\mathrm{Conf}_n(S^1, \{0\})\big) = \begin{cases} \mathbb{Z} & i = 0 \\ \mathbb{Z}^n & i = 1 \\ 0 & \textit{else} \end{cases}$$

Proof. This proof is a slightly extended version of the one in [CL16]. We compute the homology one example at a time.

First: The combinatorial model of this space is a disjoint union of zero-dimensional cubes because there is no essential vertex, so there cannot be any higher-dimensional

cell. The model has one cube for each ordering of the particles, describing the bijection to Σ_n.

Second: Remember that in this special case we think of S^1 consisting of one vertex and one edge (and we call that vertex essential even though it has valence 2). The combinatorial model, therefore, consists of zero and one-dimensional cubes. Each of the vertices of the combinatorial model has valence two, and for each cyclic ordering of the particles, there is one copy of S^1. Moving along one of the copies of S^1 corresponds to moving one particle from one end of the edge to the other end via the vertex, one particle after the other until the initial configuration is restored.

Third: The interval with one sink has contractible configuration space: we can just pull all particles into the sink.

Fourth: The interval with two sinks has connected configuration spaces by pulling particles onto one of the sinks. The combinatorial model is one-dimensional, so we only need to compute the Euler characteristic. There is a zero cube for every distribution of particles onto the two sinks, which means that there are 2^n of them. We have a 1-cell for each choice of one moving particle and every distribution of the remaining ones onto the two sinks, so there are $n2^{n-1}$ many 1-cells. Thus, the Euler characteristic is $(2-n)2^{n-1}$, which determines the rank of the first homology group.

Fifth: The configuration space of the circle with one sink is again connected by pulling the particles onto the sink. It remains to compute the Euler characteristic of the 1-dimensional model. There is precisely one zero cell, namely the one where all particles are on the sink. There is one 1-cell for each choice of one particle moving along the edge, giving n 1-cells and therefore the Euler characteristic $1-n$. □

The last two examples will play an important role later, so we describe their homology more concretely.

From the description of the combinatorial model for $\mathrm{Conf}_n(S^1, \{0\})$ it is clear that its homology is generated by one particle moving along the edge and the other particles staying fixed on the sink.

For the interval with two sinks we first consider the cases of few particles: with only one particle there is no one-dimensional class, and with two particles there is precisely one 1-class: both particles sit on the first sink, particle 1 moves to the second sink, particle 2 follows, particle 1 returns to the first sink and finally also particle 2 moves back to the first sink, see Figure 1.4.

The next proposition shows that this is the only class we need to understand, all classes in the case of three or more particles can be written as sums of those classes involving only two particles.

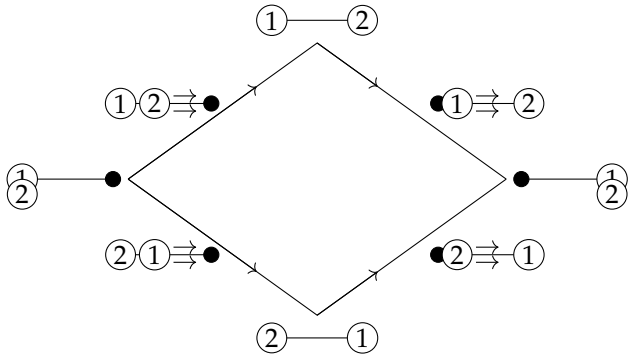

Figure 1.4: The combinatorial model of $\mathrm{Conf}_2(I, \{0,1\})$ consists of four edges.

Proposition 1.24. *The map*

$$\bigoplus_{s \neq s' \in \underline{n}} \bigoplus_{\phi} H_1(\mathrm{Conf}_{\{s,s'\}}(I,\{0,1\})) \to H_1(\mathrm{Conf}_n(I,\{0,1\}))$$

is surjective. The second indexing set is given by all maps $\phi\colon \underline{n} - \{s, s'\} \to \{0,1\}$ *and the maps out of the direct summands are induced by putting the remaining particles onto the sinks according to* ϕ*.*

Proof. Let Z be a connected cellular 1-cycle in the one-dimensional combinatorial model of $\mathrm{Conf}_n(I,\{0,1\})$ (meaning that the union of all 1-cubes with non-trivial coefficient is connected) and let $k \in \underline{n}$ be a particle. Let ξ_k and $\xi_{k'}$ be 1-cubes moving particles $k \neq k'$ along the interval such that they intersect in a single 0-cube. On both 1-cubes, the positions of the particles $\underline{n} - \{k, k'\}$ on the sinks are the same. Using the corresponding map

$$H_1(\mathrm{Conf}_{\{k,k'\}}(I,\{0,1\})) \to H_1(\mathrm{Conf}_n(I,\{0,1\}))$$

the standard generator of $H_1(\mathrm{Conf}_{\{k_1,k_2\}}(I,\{0,1\}))$ maps to a class represented by $\pm(\xi_k + \xi_{k'} - \overline{\xi_k} - \overline{\xi_{k'}})$, where $\overline{\xi_k}$ is given by the 1-cube ξ_k with k' on the other sink, and analogously for $\overline{\xi_{k'}}$. Using this cycle we can replace the 1-cubes ξ_k and $\xi_{k'}$ by $\overline{\xi_k}$ and $\overline{\xi_{k'}}$, which simply changes the order in which the particles k and k' move along the interval. Notice that this does not change the number of summands of Z. If this process disconnects Z, we look at the connected components individually.

Repeating this, we can arrange that Z is a connected cycle such that the union Z_k of all edges of Z moving k from one sink to the other is connected. The positions of the particles $\underline{n} - \{k\}$ are fixed in Z_k and the particle k does not move in $Z - Z_k$. Therefore, Z_k itself is a cycle and hence in the image of a map

$$0 = H_1(\mathrm{Conf}_{\{k\}}(I,\{0,1\})) \to H_1(\mathrm{Conf}_n(I,\{0,1\})),$$

which shows that $[Z_k] = 0$. We are then left with a cycle $Z - Z_k$ in which k is fixed.

Repeating this argument for all particles $k \in \underline{\mathbf{n}}$ eventually leads to a trivial cycle by induction on the number of summands of Z. □

Remark 1.25. The images of these different maps have lots of relations, and if one is interested in a concrete basis of the homology one can construct one in the following way: the combinatorial model of $\mathrm{Conf}_n(I, \{0,1\})$ is a graph, so it is sufficient to describe a maximal tree in this model to give a basis of the fundamental group and the first homology. As maximal tree one can take all edges where the label of the moving particle is smaller than the labels of all particles sitting on the second sink. The basis for the homology is then given by all edges where the label of the moving particle x is bigger than the label of at least one of the particles on the second sink. The corresponding cycle moves x from the first to the second sink, moves the particle with the smallest label from the second to the first sink until x is again on the first sink and finally moves those particles back to the second sink (in decreasing order). △

Another case which we can compute explicitly is the case of star graphs, or more generally graphs with exactly one essential vertex.

Proposition 1.26 ([Lü14, Proposition 3.5, p. 36])**.** *Define for* $k, \ell \in \mathbb{N}$ *the graph* Y_k^ℓ *having one vertex* v*,* k *leaves and* ℓ *edges forming self-loops. If* $k + 2\ell \geqslant 3$ *we have*

$$H_i\big(\mathrm{Conf}_n(Y_k^\ell, \varnothing)\big) = \begin{cases} \mathbb{Z} & i = 0 \\ \mathbb{Z} + (-\chi(\mathrm{UConf}_n(Y_k^\ell))) \cdot [\mathbb{Z}\Sigma_n] & i = 1 \\ = \mathbb{Z}^{1 + \frac{(n+k+\ell-2)!}{(k+\ell-1)!}(n(k+2\ell-2)-(k+\ell-1))} & \\ 0 & \textit{else.} \end{cases}$$

The proof works by calculating the Euler characteristic of the combinatorial model: the combinatorial model is 1-dimensional and the condition $k + 2\ell \geqslant 3$ implies that the configuration space is connected.

In the case $n = 2, k = 3$ and $\ell = 0$ we have two particles in the Y-shaped graph and the formula above implies that the first homology group is one-dimensional. A generating cycle is easy to visualize: start with the two particles on different leaves. Choose one of the particles and move it to the empty leaf. Now move the other particle to the leaf that just became empty and repeat this procedure until the initial configuration is obtained again, see Figure 1.5.

Remark 1.27. These calculations imply that if we have enough particles, then the dimension of the combinatorial model of $\mathrm{Conf}_n(G, Z)$ is equal to the homological dimension of the space. Indeed, choose two particles for each essential non-sink vertex and each edge incident to two sinks. Embed Y-shaped graphs without sinks into the stars of these essential vertices and copies of $(I, \{0,1\})$ onto the edges

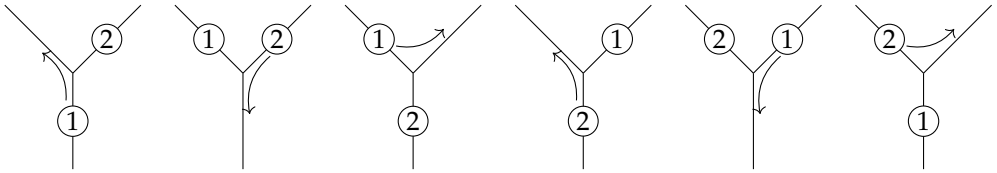

Figure 1.5: A generator of the first homology of the configuration space of two particles in the Y graph. The arrows indicate two edges of the combinatorial model, so the whole cycle is a sum of 12 edges.

connecting sinks. $H_1(\mathrm{Conf}_2(Y, \varnothing)$ and $H_1(\mathrm{Conf}_2(I, \{0, 1\}))$ are 1-dimensional, so we can choose the 1-cycles described above in each of those configuration spaces. These are embedded circles in the respective combinatorial models.

Now take the product cycle (see Definition 1.15) in the combinatorial model of $\mathrm{Conf}_n(G, Z)$, which is a torus of top-dimensional cells, see Figure 1.6. This is well-defined since the non-sink vertices and edges incident to two sinks are distinct, so we can realize any combination of movements simultaneously. This is a cycle, and since there are no higher-dimensional cells this does not represent the zero homology class. Therefore, the dimension of the combinatorial model is precisely the homological dimension, and thus our combinatorial model, in that case, has the smallest dimension possible. △

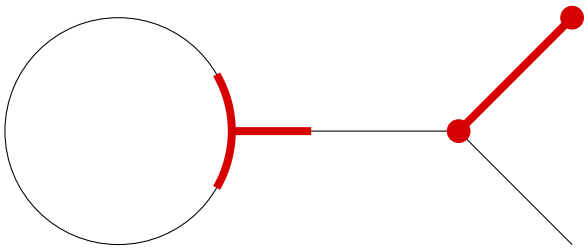

Figure 1.6: A top-dimensional homology class in the configuration space of 4 particles in this graph G. Take the standard 1-class of the particles 1 and 2 in the star and the standard 1-class of the particles 3 and 4 in the interval, then take the product of these two circles. This gives a torus, where the horizontal and vertical directions on the torus correspond to the movement of the particles in the star graph and the interval, respectively.

Chapter 2

The Mayer-Vietoris spectral sequence for configuration spaces

In this chapter we introduce the main technical tool apart from sinks in our approach to computing the homology of configuration spaces of graphs: the Mayer-Vietoris spectral sequence. We first describe the spectral sequence for general open covers of spaces and then describe how we create such a cover of $\mathrm{Conf}_n(X, Z)$ from an open cover of X. Then, we describe our main strategy of computing the E^2-page via comparison of these spectral sequences for different graphs.

2.1 The Mayer-Vietoris spectral sequence

In this section we define the Mayer-Vietoris spectral sequence for arbitrary open covers of a topological space X and prove that it converges to the homology of X. This spectral sequence is well-known, but since we could not find a source with a concise, self-contained proof, we provide one here.

Proposition 2.1. *Let X be a topological space, let J be a totally ordered (possibly uncountable) index set and let $\mathcal{U} = \{U_j\}_{j\in J}$ be an open cover of X. Writing $U_{j_0\cdots j_p} = U_{j_0} \cap \cdots \cap U_{j_p}$, there is a spectral sequence*

$$ {}^{\mathcal{U}}E^1_{p,q} = \bigoplus_{j_0<\cdots<j_p} H_q(U_{j_0\cdots j_p}) \Rightarrow H_*(X) $$

converging to the singular homology of X. This spectral sequence is called Mayer-Vietoris spectral sequence. *The boundary map $d_1 \colon E^1_{p,q} \to E^1_{p-1,q}$ is given by the alternating sum*

$$ d_1 = \Sigma_{i=0}^{p} (-1)^i \delta_i, $$

where

$$\delta_i \colon \bigoplus_{j_0<\cdots<j_p} H_q(U_{j_0\cdots j_p}) \to \bigoplus_{j_0<\cdots<j_{p-1}} H_q(U_{j_0\cdots j_{p-1}})$$

is induced by the inclusion maps

$$U_{j_0} \cap \cdots \cap U_{j_p} \hookrightarrow U_{j_0} \cap \cdots \cap \widehat{U}_{j_i} \cap \cdots \cap U_{j_p},$$

where the hat indicates that U_{j_i} *is removed from the intersection.*

In the proof we need the following standard result.

Theorem 2.2 ([McC01, Thm 2.15, p. 48])**.** *Let* $\{M^{*,*}, d', d''\}$ *be a first quadrant double complex. There are two associated spectral sequences associated with* M*, one with boundary maps* $(d_0, d_1) = (d', d'')$ *and one with* $(d_0, d_1) = (d'', d')$*, and both converge to the homology of the total complex of* M*.*

Proof of Proposition 2.1. Define

$${}^{\mathcal{U}}E^0_{p,q} = E^0_{p,q} := \bigoplus_{j_0<\cdots<j_p} C^{\mathrm{sing}}_q(U_{j_0\cdots j_p}),$$

where C^{sing}_q denotes the abelian group of singular q-chains. Let

$$d'_{p,q} \colon E^0_{p,q} \to E^0_{p,q-1}$$

be the direct sum of the differentials of the singular chain complexes

$$C^{\mathrm{sing}}_q(U_{j_0\cdots j_p}) \to C^{\mathrm{sing}}_{q-1}(U_{j_0\cdots j_p}),$$

multiplied with the sign $(-1)^{p+1}$, and let

$$d''_{p,q} \colon E^0_{p,q} \to E^0_{p-1,q}$$

be given by $d'' = \Sigma_{i=0}^{p}(-1)^i\delta_i$, where

$$\delta_i \colon \bigoplus_{j_0<\cdots<j_p} C^{\mathrm{sing}}_q(U_{j_0\cdots j_p}) \to \bigoplus_{j_0<\cdots<j_{p-1}} C^{\mathrm{sing}}_q(U_{j_0\cdots j_{p-1}})$$

is induced by the inclusion maps (see the definition of d_1).

It is straightforward to check that this is, in fact, a double complex, i.e. that $d'^2 = 0$, $d''^2 = 0$ and $d'd'' = -d''d'$. Therefore, by Theorem 2.2, this determines two spectral sequences converging to the same homology.

$\underline{d_0 = d'}$: If we first take homology in the d'-direction, then the first page will be given by

$$E^1_{p,q} = \bigoplus_{j_0<\cdots<j_p} H_q(U_{j_0\cdots j_p}),$$

and the d_1 differential will be induced by d''. This is precisely the spectral sequence we described in the statement, so it remains to show that the homology it converges to is actually the homology of X. We will do this by computing the other spectral sequence.

$\underline{d_0 = d''}$: If we first take homology in the d''-direction, then we claim that the infinity page only consists of one single column, namely the zeroth column.

Let $C_\bullet^{\mathcal{U}}(X) \subset C_\bullet^{\text{sing}}(U)$ be the chain complex of $\mathcal{U}$-small singular simplices, i.e. singular simplices whose image is contained in one of the open sets $U \in \mathcal{U}$. Given a $\mathcal{U}$-small simplex $\sigma\colon \Delta^q \to X$ let $\mathcal{U}_\sigma$ be defined as

$$\mathcal{U}_\sigma := \{U \in \mathcal{U} \mid \sigma \subset U\}.$$

The realization of the nerve of $\mathcal{U}_\sigma$ is a (possibly infinite) simplex because the intersection of all open sets is non-trivial (it contains at least σ), so we denote it by $\Delta_{\mathcal{U}_\sigma}$. If $\mathbb{Z}_\sigma$ denotes the free $\mathbb{Z}$-module on the single generator σ, then we have the canonical inclusion of chain complexes of the *cellular* chain complex of $\Delta_{\mathcal{U}_\sigma}$ with coefficients in $\mathbb{Z}_\sigma$ into the q-th row of the E^0-page:

$$\begin{array}{ccccccc}
\cdots \leftarrow & \bigoplus_{j_0<\cdots<j_{p-1}} C_q^{\text{sing}}(U_{j_0} \cap \cdots \cap U_{j_{p-1}}) & \leftarrow & \bigoplus_{j_0<\cdots<j_p} C_q^{\text{sing}}(U_{j_0} \cap \cdots \cap U_{j_p}) & \leftarrow \cdots \\
& \uparrow & & \uparrow & \\
\cdots \leftarrow & C_{p-1}^{\text{cell}}(\Delta_{\mathcal{U}_\sigma}; \mathbb{Z}_\sigma) & \leftarrow & C_p^{\text{cell}}(\Delta_{\mathcal{U}_\sigma}; \mathbb{Z}_\sigma) & \leftarrow \cdots
\end{array}$$

Summing over all $\mathcal{U}$-small singular q-simplices σ in X, it is straightforward to check that this actually gives an isomorphism of chain complexes:

$$\bigoplus_\sigma C_\bullet^{\text{cell}}(\Delta_{\mathcal{U}_\sigma}; \mathbb{Z}_\sigma) \xrightarrow{\cong} E^0_{\bullet,q}.$$

The homology of the left-hand chain complex is concentrated in degree zero, where it is given by $C_q^{\mathcal{U}}(X)$, so the first page of the spectral sequence is concentrated in the zeroth column:

$$E^1_{p,q} = \begin{cases} C_q^{\mathcal{U}}(X) & \text{if } p = 0, \\ 0 & \text{else.} \end{cases}$$

The vertical differentials are induced by d' and therefore given by the ordinary singular chain complex boundary maps. By [Hat02, Proposition 2.21, p.119], $\mathcal{U}$-small singular simplices are enough to compute the homology of X, i.e. the inclusion

$$C_\bullet^{\mathcal{U}}(X) \hookrightarrow C_\bullet^{\text{sing}}(X)$$

induces an isomorphism on homology. Therefore, the E^2-page is equal to the E^∞-page and given by

$$E^1_{p,q} = \begin{cases} H^{\mathrm{sing}}_q(X) & \text{if } p = 0, \\ 0 & \text{else,} \end{cases}$$

and thus the Mayer-Vietoris spectral sequence as defined in the statement of the proposition converges to the homology of X. □

2.1.1 Higher boundary maps

For concrete calculations later in this thesis, we now give a description of the higher boundary maps d_k for $k \geqslant 2$.

We first describe the map

$$d_2 \colon E^2_{p,q} \to E^2_{p-2,q+1}.$$

Let $Z \in E^0_{p,q}$ represent an element $[Z] \in E^2_{p,q}$, meaning $d_0(Z) = 0 \in E^0_{p,q-1}$ and $[d_1(Z)] = 0 \in E^1_{p-1,q}$. By this latter equality, there exists an element $Z_2 \in E^0_{p-1,q+1}$ such that $d_0(Z_2) = d_1(Z)$. The element $d_1(Z_2) \in E^0_{p-2,q+1}$ then represents $d_2([Z]) \in E^2_{p-2,q+1}$, see Figure 2.1. A simple diagram chase shows that this definition of $d_2([Z])$ is independent of the choice of preimage of $d_1(Z)$ under d_0. We will not make a notational distinction between the maps d_k on the E^k-page and their corresponding maps on the zeroth page (which are well-defined only up to choices).

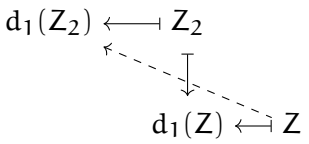

Figure 2.1: Constructing the boundary map d_2.

The choice of Z_2 can be interpreted as follows: the element $d_1(Z)$ removes open sets from the $(p+1)$-fold intersections and maps the cycles into those bigger open sets. Considered as cycles in a bigger space, some of them might now be hit by a boundary map and therefore represent zero, see Figure 2.2. Choosing a preimage Z_2 of $d_1(Z)$ under d_0 is precisely choosing explicit singular chains bounding the null-homologous cycles of $d_1(Z)$.

If we now have $Z \in E^0_{p,q}$ representing an element $[Z] \in E^3_{p,q}$, then we choose a representative $d_1(Z_2)$ of $d_2([Z])$ as above. Since $[Z]$ is an element of the third page, we have $d_2([Z]) = 0$ on the second page, meaning that

$$[d_1(Z_2 + Z')] = 0 \in E^1_{p-2,q+1}$$

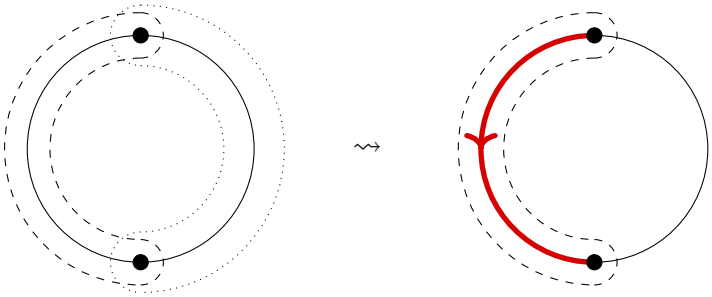

Figure 2.2: The left picture depicts the difference of two 0-cycles in the intersection of two open subsets of S^1. The 0-cycle on the top has sign -1, the other one $+1$. Removing an open set from the intersection makes these cycles homologous, as visible on the picture on the right. The red 1-cell bounds the two 0-cycles in this contractible space.

for some element $Z' \in E^0_{p-1,q+1}$ representing an element of $E^1_{p-1,q+1}$. Now choose an element $Z_3 \in E^0_{p-2,q+2}$ such that

$$d_0(Z_3) = d_1(Z_2 + Z').$$

Then $d_1(Z_3)$ represents $d_3([Z]) \in E^3_{p-3,q+2}$. For higher boundary maps, the process is the same: add images of all lower boundary maps until the element represents zero on the first page, and then choose explicit chains bounding these null-homologous cycles. Chasing diagrams one again shows that these choices do not affect the final result.

Often, the addition of images of lower boundary maps (the Z' in the description of d_3) is not necessary. Then, computing the boundary maps proceeds by repeated application of the following two steps:

- forget open sets from the intersections, and
- choose explicit chains bounding the resulting cycles.

2.1.2 Interpreting classes on the E^∞-page

Once we computed the ∞-page, we have to interpret the classes of $E^\infty_{p,q}$ as classes in $H_{p+q}(X)$. In this section we explain how this works.

Let $Z \in E^0_{p,q}$ represent an element of $E^\infty_{p,q}$, then we know that $d_k([Z]) = 0 \in E^k_{p-k,q+k-1}$ for all $k \geqslant 0$. In particular, this is true for $k = p$, and we get by the procedure described above an element $Z_p \in E^0_{0,p+q}$ such that $d_0(Z_p)$ represents $d_p([Z]) \in E^p_{0,p+q-1}$. Mapping this element Z_p via

$$E^0_{0,p+q} = \bigoplus_{U \in \mathcal{U}} C_{p+q}(U) \to C_{p+q}(X)$$

yields a cycle in X that represents the homology class in $H_{p+q}(X)$ corresponding to $[Z] \in E^{\infty}_{p,q}$, see Figure 2.3.

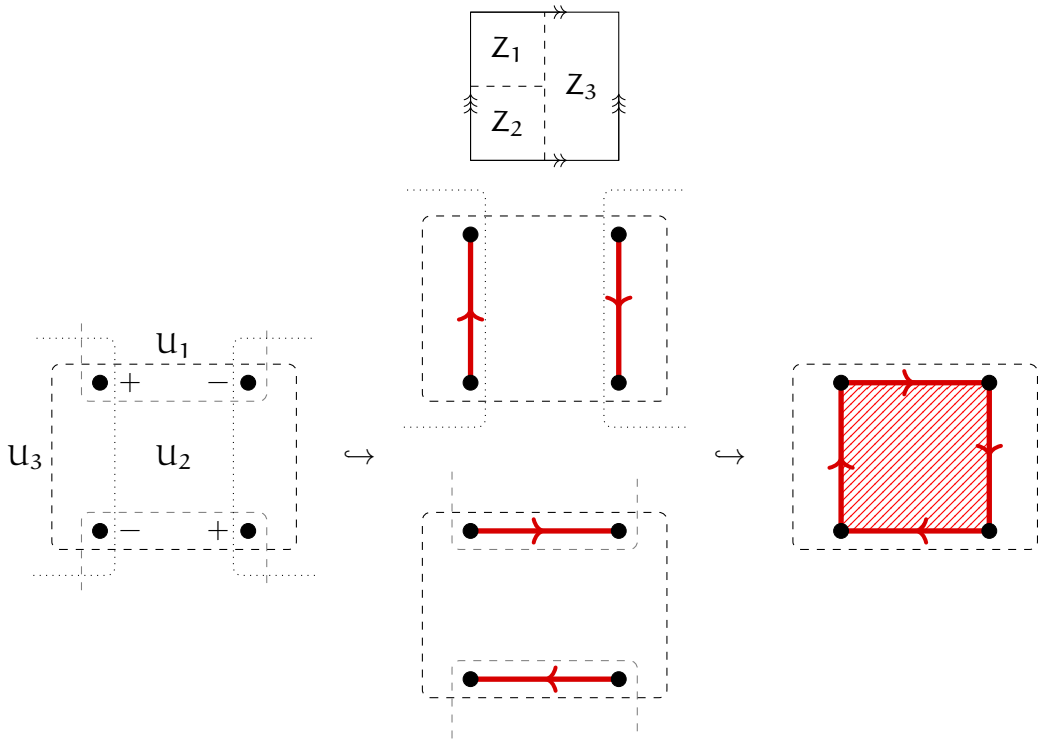

Figure 2.3: Interpreting elements on the E^{∞}-page. Consider the Mayer-Vietoris spectral sequence for the open cover of the torus given by U_i a small neighborhood of Z_i for $1 \leqslant i \leqslant 3$ as in the picture on the top. The fundamental class of the torus is represented by the depicted linear combination of zero-cycles in $U_1 \cap U_2 \cap U_3$, i.e. as an element of $E^{\infty}_{2,0}$. Lifting this element to the zeroth column is done by forgetting open sets from the intersection, one at a time. In each of the three open sets this produces a 2-cube as visualized above, and mapped to the total space they glue together to give the whole torus.

2.2 A Mayer-Vietoris spectral sequence for configuration spaces

In this section we will recall the Mayer-Vietoris spectral sequence for *configuration spaces* as constructed and studied in [CL16]. This, together with the notion of configuration spaces with sinks, is the main technical tool through which we understand the homology of configuration spaces of graphs.

Definition 2.3 ([CL16, Definition 2.7, p. 8])**.** Let (X, Z) be a space with sinks and let $\mathcal{V} = \{V_i\}_{i \in I_\mathcal{V}}$ be a totally ordered open cover of X. Then $\mathcal{V}^S$ is the following open cover of $\mathrm{Conf}_S(X, Z)$: there is one open set U_ϕ for each map

$$\phi \colon S \to I_\mathcal{V}$$

given by

$$U_\phi := \bigcap_{s \in S} \pi_s^{-1}(V_{\phi(s)}),$$

where π_s is the projection to the position of the particle s, see Definition 1.10. We have

$$\begin{aligned} U_{\phi_0 \cdots \phi_p} &:= U_{\phi_0} \cap \cdots \cap U_{\phi_p} \\ &= \bigcap_{j=0}^{p} \bigcap_{s \in S} \pi_s^{-1}\left(V_{\phi_j(s)}\right) \\ &= \bigcap_{s \in S} \pi_s^{-1}\left(\bigcap_{j=0}^{p} V_{\phi_j(s)}\right) \\ &= \bigcap_{s \in S} \pi_s^{-1}\left(V_{\phi_0(s) \cdots \phi_p(s)}\right). \end{aligned}$$

This is a Σ_S-equivariant open cover of $\mathrm{Conf}_S(X, Z)$, meaning that for each $U \in \mathcal{V}^S$ and permutation $\eta \in \Sigma_S$ the open set $U \cdot \eta$ is also in $\mathcal{V}^S$. Choosing an arbitrary ordering of S, order the maps ϕ lexicographically.

The *Mayer-Vietoris spectral sequence* ${}^{\mathcal{V}^S}E^*_{\bullet,\bullet}$ *for* $\mathrm{Conf}_S(X, Z)$ *associated with the open cover* $\mathcal{V}$ *of* X is defined to be the Mayer-Vietoris spectral sequence for $\mathcal{V}^S$ as defined in Proposition 2.1. △

By the description above, the E^1-page is given by

$$E^1_{p,q} = \bigoplus_{\phi_0 < \cdots < \phi_p} H_q\left(\bigcap_{s \in S} \pi_s^{-1}\left(V_{\phi_0(s) \cdots \phi_p(s)}\right)\right).$$

The boundary map can be thought of as weakening the restrictions put onto the individual particles.

Remark 2.4. Let $\mathcal{V}_1$ and $\mathcal{V}_2$ be ordered open covers of spaces X and Y, respectively. If there is a map $f \colon X \to Y$ such that for each $V \in \mathcal{V}_1$ the image $f(V)$ is contained completely in precisely one open set in $\mathcal{V}_2$, then we get an induced map $f_* \colon \mathcal{V}_1 \to \mathcal{V}_2$. If this map is order-preserving, then the induced map $\mathcal{V}_1^S \to \mathcal{V}_2^S$ is for each S order-preserving since we order lexicographically.

Therefore, this induces a map of the corresponding Mayer-Vietoris spectral sequences as above: On the zero page it is induced by the map f restricted to the open sets

$$U_{\phi_0 \cdots \phi_p} \to U_{f_*\phi_0 \cdots f_*\phi_p}$$

if all $f_*\phi_i$ are distinct, and the zero map else. In this way, the Mayer-Vietoris spectral sequence is functorial with respect to such maps f. △

We will now describe this spectral sequence for the case where (X, Z) is a graph with sinks and $\mathcal{V}$ is a very special kind of open cover. For the rest of this section, we will rename X to G. Furthermore, we assume that Z is a union of edges and vertices of G.

Proposition 2.5. *Let S be a finite set and let $\mathcal{V}$ be an open cover of G such that*

- *for each $V \in \mathcal{V}$ the space $\mathrm{Conf}_n(V, Z_V)$ for $Z_V := Z \cap V$ is connected for all n, and*
- *the intersection of two or more open sets is a disjoint union of star graphs and intervals, such that each essential vertex is a sink and every edge between two sinks is itself a sink.*

Then each intersection

$$U_{\phi_0} \cap \cdots \cap U_{\phi_p}$$

of open sets $U_{\phi_j} \in \mathcal{V}^S$ is homotopy equivalent to a disjoint union of spaces of the form

$$\prod_{i \in I_{\mathcal{V}}} \mathrm{Conf}_{S_i}(V_i, Z_{V_i})$$

for $\sqcup_{i \in I_{\mathcal{V}}} S_i \subset S$.

Corollary 2.6. *If additionally $I_{\mathcal{V}}$ is finite and the configuration spaces of all (or all but one) V_i have free homology, then the q-th homology of such a space as above is given by*

$$\bigoplus_{\Sigma_{i \in I_{\mathcal{V}}} q_i = q} \bigotimes_{i \in I_{\mathcal{V}}} H_{q_i}(\mathrm{Conf}_{S_i}(V_i, Z_{V_i}))$$

and the module $E^1_{p,q}$ of the corresponding Mayer-Vietoris spectral sequence is given by

$$E^1_{p,q} \cong \bigoplus_{j \in J_{pq}} \bigoplus_{\Sigma_{i \in I_{\mathcal{V}}} q_i = q} \bigotimes_{i \in I_{\mathcal{V}}} H_{q_i}(\mathrm{Conf}_{S_i^j}(V_i, Z_{V_i}))$$

for some index set J_{pq} and some tuples of sets $(S^j_\bullet)$ such that $\sqcup_{i \in I_{\mathcal{V}}} S^j_i \subset S$.

Remark 2.7. It would require some work to give the index set J_{pq} explicitly: for each intersection $U_{\phi_0} \cap \cdots \cap U_{\phi_p}$, there are a lot of connected components determining different distributions of the particles into the intersections. A particle is restricted to the intersection of two or more open sets either because different ϕ_j map it to different indices or because it is blocked by other particles. In particular, a particle can be restricted to an intersection in one connected component of such an intersection while being able to move freely in one of the $V \in \mathcal{V}$ in another. △

Proof of Proposition 2.5. In $U_{\phi_0 \cdots \phi_p}$ each particle is restricted to a certain intersection of the open sets $V_i \in \mathcal{V}$, namely $V_{\phi_0(i) \cdots \phi_p(i)}$. Let X be one path component of $U_{\phi_0 \cdots \phi_p}$ and let F_I be the set of particles which are always in the interior of some edge for all configurations of X. This might happen either because the particles themselves are restricted by the maps ϕ_j or because they are blocked by another particle, see the remark preceding this proof. Up to homotopy, these particles do not move at all, so the projection map π_{S-F_I} induces a homotopy equivalence $X \simeq \pi_{S-F_I}(X)$.

If in $\pi_{S-F_I}(X)$ there are still particles restricted to an intersection of two or more open sets, then these intersections have to be stars with sinks by the properties of the open cover. Therefore, we can pull all these particles F_{star} onto their corresponding sinks, showing that the map $\pi_{S-F_I-F_{star}}$ induces a homotopy equivalence

$$X \simeq \pi_{S-F_I-F_{star}}(X) =: X'.$$

In X' every particle is free to move inside one of the open sets V_i, and we denote the set of particles by $S' := S - F_I - F_{star}$. The particles in V_i will be denoted by $S_i \subset S'$.

We will now deform our open sets V_i such that their intersections are contained in Z. For each edge e incident to a sink vertex v, there is at most one open set V_i containing both vertices of e by the properties listed in the statement. The intersection of e with each of the other open sets of $\mathcal{V}$ containing v can be collapsed to the sink by sliding the particles onto v, see Figure 2.4. Repeating this for all edges incident to sink vertices we can arrange that the interior of each such edge intersects at most one of these deformed sets.

If two or more of the deformed open sets now intersect outside of sinks, then this intersection is on an edge e between non-sink vertices. This can only be the intersection of precisely two open sets, each containing one of the two distinct vertices because otherwise, an intersection of two of the open sets would contain a non-sink vertex. Now shrink the two open sets on that edge e until they are disjoint, pulling the particles along, see Figure 2.5.

The resulting collection of sets $\{V_i'\}$ has the property that the intersection of any two of them is contained in Z, so the particles in V_i' and V_j' for $i \neq j$ do not interact with each other. This identifies X' with a connected component of

$$\mathrm{Conf}_{S'}\left(\sqcup_i V_i', \sqcup_i Z_{V_i'}\right),$$

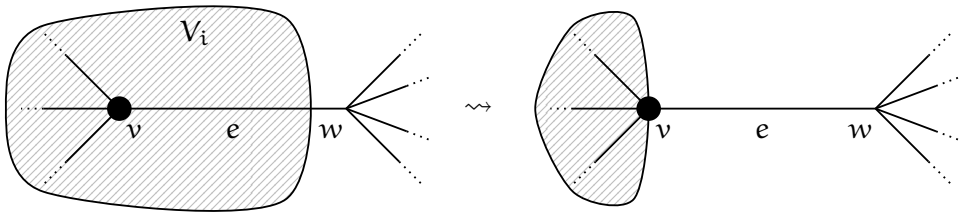

Figure 2.4: Shrinking one of the V_i containing only a part of the edge e. The right part of the picture shows the deformed set, which is disjoint from the interior of e.

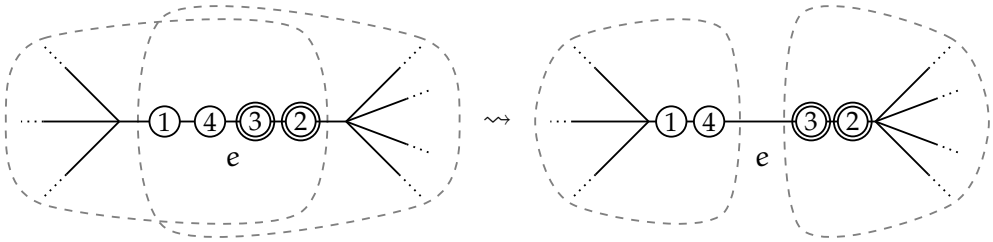

Figure 2.5: Shrinking two of the deformed V_i intersecting in an edge e. The particles 1 and 4 are restricted to the left set, the other two particles are restricted to the right set. The right part of the picture shows the deformed sets, where the two sets do not intersect inside e anymore.

namely we have

$$X' \simeq \prod_i \mathrm{Conf}_{S_i}(V_i', Z_{V_i'}).$$

Since the deformation above is, in fact, a homotopy equivalence as spaces with sinks $(V_i, Z_{V_i}) \simeq (V_i', Z_{V_i'})$, this shows by Corollary 1.8 that

$$X \simeq X' \simeq \prod_i \mathrm{Conf}_{S_i}(V_i', Z_{V_i'}) \simeq \prod_i \mathrm{Conf}_{S_i}(V_i, Z_{V_i}). \qquad \square$$

Remark 2.8. By the result above, we can investigate the spectral sequence by understanding the configuration spaces of the open sets V_i. By Proposition 1.22, we can replace these open sets by closed sets, which turns them into finite graphs again. $\triangle$

2.3 The sink comparison argument

In this section we will always assume that (G, Z) is a graph with sinks and Z is a union of edges and vertices of G.

The number of open sets and their intersections is quite large for the cover $\mathcal{V}^S$ as described above, which means that the combinatorics of the corresponding Mayer-Vietoris spectral sequences are usually quite involved. For this reason, we will investigate them mostly by comparing the E^1- and E^2-pages of these spectral sequences for related configuration spaces. We will describe the general idea for such comparisons in this section.

Assume we are in a situation as in Corollary 2.6 with $\mathcal{V} = \{V_1, V_2\}$ and look at the first page of the spectral sequence:

$$\begin{aligned} E^1_{p,q} &\cong \bigoplus_{j\in J} \bigoplus_{q_1+q_2=q} H_{q_1}(\mathrm{Conf}_{S_1^j}(V_1, Z_{V_1})) \otimes H_{q_2}(\mathrm{Conf}_{S_2^j}(V_2, Z_{V_2})) \\ &\cong \bigoplus_{q_1+q_2=q} \bigoplus_{j\in J} H_{q_1}(\mathrm{Conf}_{S_1^j}(V_1, Z_{V_1})) \otimes H_{q_2}(\mathrm{Conf}_{S_2^j}(V_2, Z_{V_2})). \end{aligned}$$

The horizontal boundary map d_1 is induced by removing open sets from the intersections. In particular, the map d_1 does not change the degree of the homology groups, so the rows $E^1_{\bullet,q}$ as above split into a direct sum of chain complexes $\mathcal{C}_{q_1,q_2}$ with modules of the form

$$\bigoplus_{j\in J} H_{q_1}(\mathrm{Conf}_{S_1^j}(V_1, Z_{V_1})) \otimes H_{q_2}(\mathrm{Conf}_{S_2^j}(V_2, Z_{V_2})).$$

for each fixed choice of q_1 and q_2 satisfying $q_1 + q_2 = q$.

Now assume that $q_1 = 0$, then we have

$$H_0(\mathrm{Conf}_{S_1^j}(V_1, Z_{V_1})) = \mathbb{Z}$$

for all j by the choice of our open cover. If $\widetilde{Z}_{V_1} \supset Z_{V_1}$ is a bigger set of sinks, then the inclusion of graphs with sinks induces an isomorphism on zeroth homology:

$$H_0(\mathrm{Conf}_{S_1^j}(V_1, Z_{V_1})) \xrightarrow{\cong} H_0(\mathrm{Conf}_{S_1^j}(V_1, \widetilde{Z}_{V_1})).$$

In particular, the chain complex $\mathcal{C}_{0,q_2}$ is also part of the q_2-th row of the E^1-page of $\mathrm{Conf}_S(G, Z \cup \widetilde{Z}_{V_1})$. It turns out that in certain situations one can choose the sinks in such a way that the q_2-th row is actually *equal* to $\mathcal{C}_{0,q_2}$. If we can compute the homology of this row by knowledge about the space with more sinks, then we get a part of the q_2-th row of the E^2-page of our original spectral sequence.

Example 2.1. Take the H-graph given by gluing together two stars of valence 3 along one of their leaves. Give H the path metric such that each internal edge has length 1 and each leaf has length $1/2$. Consider the open cover V_v, V_w given by the open balls of radius 1 around the two vertices v and w. Their intersection is the interior of the horizontal edge of H, so we are in the situation of Corollary 2.6.

In the **zeroth row** all entries will be direct sums of modules of the form

$$H_0(\mathrm{Conf}_{S_v}(V_v)) \otimes H_0(\mathrm{Conf}_{S_w}(V_w)) \cong \mathbb{Z} \otimes \mathbb{Z}$$

for different choices of $S_v \sqcup S_w \subset S$. By the arguments above, the zeroth row of the first page of the analogous spectral sequence for $\mathrm{Conf}_S(H, \{v, w\})$ is exactly the same. On the first page of that latter spectral sequence, however, all q-th rows for $q > 0$ are zero. Therefore, the second page is already the infinity page and the zeroth row of the E^2-pages of *both* spectral sequences is given by $H_\bullet(\mathrm{Conf}_S(H, \{v, w\}))$. This is, in fact, the same as $H_\bullet([0,1], \{0,1\})$ by pulling the particles from the leaves to the sinks. By Proposition 1.23 we know that this homology is concentrated in the zeroth and first degree.

The **first row** of the original spectral sequence splits into two chain complexes, one with all modules of the form

$$H_1(\mathrm{Conf}_{S_v}(V_v)) \otimes H_0(\mathrm{Conf}_{S_w}(V_w)),$$

and one with modules

$$H_0(\mathrm{Conf}_{S_v}(V_v)) \otimes H_1(\mathrm{Conf}_{S_w}(V_w)).$$

By symmetry, we only need to compute the homology of that first chain complex. This first chain complex does not change if we turn w into a sink. In fact, the chain complex is the same as the complete first row of the E^1-page of the corresponding spectral sequence for $\mathrm{Conf}_S(H, \{w\})$. The bottom row of this latter E^1-page is the same as before, so its homology is concentrated in the zeroth and first degree.

$\mathrm{Conf}_S(H, \{w\})$ is 1-dimensional, so for dimension reasons the first row of that chain complex can only have homology in the zeroth column, and therefore the first row of the E^2-page of our original spectral sequence is zero everywhere except for the zeroth column.

The **second row** can only have non-trivial homology for $p = 0$ for dimension reasons. The chain complex is given by modules of the form

$$H_1(\mathrm{Conf}_{S_v}(V_v)) \otimes H_1(\mathrm{Conf}_{S_w}(V_w)),$$

and by Proposition 1.22 this is the tensor product of the first homology groups of configuration spaces of star graphs. Therefore, $E^2_{0,2}$ is generated by products of two star classes (see Definition 1.15).

This describes the E^2-page and shows that this is already the E^∞-page, see Figure 2.6. Therefore, we were able to deduce the homology of the H-graph from the knowledge of configuration spaces of star graphs and the interval with two sinks.

2	M_{YY}		
1	M_Y		
0	$\mathbb{Z}$	M_H	
	0	1	2

Figure 2.6: The E^∞-page of the Mayer-Vietoris spectral sequence for the H-graph. M_{YY} is generated by products of star classes, M_Y is generated by star classes and M_H is generated by classes in $\mathrm{Conf}_n([0,1], \{0,1\})$.

Chapter 3

Torsion-freeness and generators

In this chapter we will investigate the homology of configuration spaces of graphs. More specifically, we show torsion-freeness in the case of trees with loops, show that the homology is generated by products of 1-classes in that case and give a description of the 1-classes for general graphs.

In the proofs of those results it will be useful to consider configuration spaces with sinks in order to compute specific Mayer-Vietoris spectral sequences. Many of our proofs will, therefore, be done in this more general setting. At the end of this chapter, we will — for the sake of completeness — then show that the results mentioned above also hold for any subset of the vertex set turned into sinks.

All of the theorems and proofs in this chapter appeared already in [CL16]. We only added a few explanations and adapted the notation.

3.1 Configurations of particles in trees with loops

In this section we will describe the homology of $\mathrm{Conf}_n(G)$ for the case where G is a tree with loops, see Definition 1.13.

Theorem A ([CL16, Theorem A, p. 2])**.** *Let G be a tree with loops and let n be a natural number. Then the integral homology $H_q(\mathrm{Conf}_n(G);\mathbb{Z})$ is torsion-free for each $q \geqslant 0$.*

The next result gives a concrete generating set for the homology of configuration spaces of such graphs, namely the set of all products of disjoint basic classes (see Definition 1.15).

Definition 3.1. For $k \geqslant 3$ let Star_k be the star graph with k leaves, let H be the tree with two vertices of valence three and let S^1 be the circle with one vertex of valence 2. We call a class $\sigma \in H_1(\mathrm{Conf}_n(G))$ *basic* if there exists a piecewise linear embedding ι of H, S^1 or Star_k for some k into G such that σ is in the image of the induced map $H_1(\mathrm{Conf}_n(\iota))$. △

Theorem B ([CL16, Theorem B, p. 2])**.** *Let* G *be a tree with loops and let* n *be a natural number. Then the homology of* $\mathrm{Conf}_n(G)$ *is generated by products of disjoint basic classes.*

Remark 3.2. Note that in such a product there may also be degree zero classes involved: the particles which are not part of any basic class need to be put somewhere on the graph. Any way of putting them disjointly from the 1-classes determines a product of disjoint basic classes. △

We will more generally prove Theorem A and Theorem B for all graphs as in the statement of the theorems with any (possibly empty) subset of the set of vertices *of valence one* turned into sinks, and the proof will proceed by induction on the number of essential vertices. The generalization of adding sinks to the vertices of valence 1 is needed in order to prove the induction step. We first prove the base case:

Proposition 3.3 ([CL16, Proposition 3.1, p. 8])**.** *Let* G *be a finite connected graph with precisely one essential vertex and* Z *a subset of the set of vertices of valence 1. Then* $H_1(\mathrm{Conf}_n(G, Z))$ *is free and generated by basic classes.*

In order to clarify the notion of a "basic class" in a graph *with sinks* we make the following remark:

Remark 3.4. Classes in $H_1(\mathrm{Conf}_n([0,1], \{0,1\}))$ can be regarded as classes in the ordinary configuration space of the H-graph $\mathrm{Conf}_n(H)$, see Figure 3.1. Replace both spaces by their combinatorial models and define a continuous map as follows: take a 0-cell of the configuration space with sinks and replace particles sitting on a sink vertex with them sitting on the corresponding lower leaf of the H-graph in their canonical ascending order. Moving a particle x from one sink vertex to the other is then given by moving all particles blocking x's path to the upper leaf, moving x onto the horizontal edge, moving the particles on the upper leaf back to the lower leaf and repeating the same game on the other side in reverse. This determines a continuous map between combinatorial models and thus induces a map on 1-cycles.

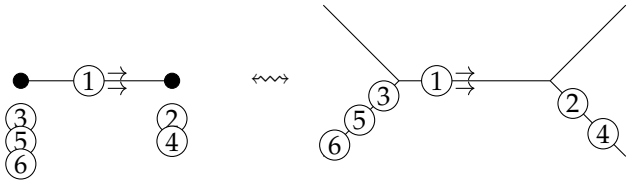

Figure 3.1: Comparing $\mathrm{Conf}_n([0,1], \{0,1\})$ and $\mathrm{Conf}_n(H)$

This map is injective in homology: composing the map with the map collapsing the two pairs of leaves to sinks gives a map that is homotopic to the identity, showing that $H_1(\mathrm{Conf}_n([0,1], \{0,1\}))$ is a direct summand of $H_1(\mathrm{Conf}_n(H))$. Compare this also to the computation of $H_*(\mathrm{Conf}_n(H))$ at the end of Section 2.3.

Because of this correspondence, the classes in $H_1(\mathrm{Conf}_n([0,1],\{0,1\}))$ are also called H-classes. When we talk about an H-class in a graph with sinks (G,Z) we hence allow the corresponding map

$$H \to G$$

to collapse some of the leaves to sinks (instead of requiring it to be an embedding as in Definition 3.1). The notion of *basic class* in the configuration space of a graph with sinks is adapted accordingly to include this new type of H-classes. △

Remark 3.5. Proposition 3.3 is indeed the full homology calculation for graphs with one essential vertex since the combinatorial model of $\mathrm{Conf}_n(G,Z)$ is 1-dimensional by Proposition 1.20, so all higher homology groups are trivial. △

In the proof we will need the following definition:

Definition 3.6. For finite sets $T \subset S$, a finite graph G, a subset $K \subset G$, and sinks $Z \subset V(G)$ write $\Gamma = (G,K)$ and define

$$\mathrm{Conf}_{S,T}(\Gamma,Z) = \{f\colon S \to G \mid f(T) \subset K\} \subset \mathrm{Conf}_S(G,Z). \qquad \triangle$$

As a consequence of the definition, we get

$$\mathrm{Conf}_{S,\varnothing}(\Gamma,Z) = \mathrm{Conf}_S(G,Z)$$

and

$$\mathrm{Conf}_{S,S}(\Gamma,Z) = \mathrm{Conf}_S(K, Z\cap K).$$

Proof of Proposition 3.3. By Proposition 1.20, $\mathrm{Conf}_n(G,Z)$ is homotopy equivalent to a graph, so the first homology is free. To see that it is generated by basic classes, we inductively use a Mayer-Vietoris long exact sequence.

For a sink $z \in Z$ define the pair of spaces $\Gamma_z = (G, G-\{z\})$. Notice that

$$\mathrm{Conf}_{S,\varnothing}(\Gamma_z,Z) = \mathrm{Conf}_S(G,Z).$$

and

$$\begin{aligned}\mathrm{Conf}_{S,S}(\Gamma_z,Z) &= \mathrm{Conf}_S(G-\{z\}, Z-\{z\})\\ &\simeq \mathrm{Conf}_S(G, Z-\{z\}),\end{aligned}$$

where the last homotopy equivalence follows from Proposition 1.22 and the fact that z has valence 1. For two sinks $z_0 \neq z_1$ we therefore have

$$\mathrm{Conf}_{S,S}(\Gamma_{z_0},Z) \simeq \mathrm{Conf}_{S,\varnothing}(\Gamma_{z_1}, Z-\{z_0\}).$$

Moving elements from $S-T$ to T and using the above identifications, we will show by induction on $|S-T|$ and the number of sinks $|Z|$ that the first homology of all

spaces $\mathrm{Conf}_{S,T}(\Gamma, Z)$ for any $\Gamma = (G, K)$ and any set of sinks Z is generated by basic classes.

In the base case we have $T = \varnothing$ and $Z = \varnothing$, so the space we are investigating is the ordinary configuration space $\mathrm{Conf}_S(G)$, which is generated by basic classes by Proposition 3.10 (this is not a circular argument, the proposition is only stated and proven later since it is the main step to compute the first homology of configuration spaces of arbitrary finite graphs). For the induction step, choose an arbitrary $s \in S - T$ and a sink $z_0 \in Z$, and take the open covering $\{U_1, U_2\}$ of $\mathrm{Conf}_{S,T}(\Gamma_{z_0}, Z)$ given by the subsets

$$U_1 := \pi_s^{-1}(G - \{z_0\}) \quad \text{and} \quad U_2 := \pi_s^{-1}(\{x \in G \mid d_G(x, z_0) < 1\}).$$

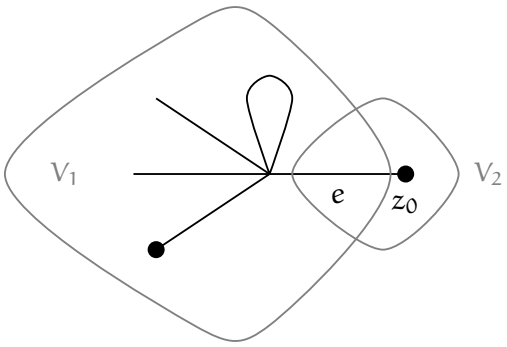

Figure 3.2: The open cover $\{U_1, U_2\}$ of the configuration space is defined by restricting particle s to one of these two open sets V_1 and V_2, respectively.

The interesting part of the Mayer-Vietoris long exact sequence is the following:

$$H_1(U_1) \oplus H_1(U_2) \to H_1(\mathrm{Conf}_{S,T}(\Gamma_{z_0}, Z)) \to H_0(U_1 \cap U_2) \\ \to H_0(U_1) \oplus H_0(U_2).$$

We have $U_1 \simeq \mathrm{Conf}_{S,T\sqcup\{s\}}(\Gamma_{z_0}, Z)$, and U_2 is homotopy equivalent to a disjoint union of the space $\mathrm{Conf}_{S-\{s\},T}(\Gamma_{z_0}, Z)$ and several copies of $\mathrm{Conf}_{S'}(G, Z - \{z_0\})$ for different finite sets $S' \subset S - \{s\}$. Those latter components of U_2 arise if particles of T sit between s and z_0, preventing s to move to the sink. The set S' is then given by the set of all particles on the other side of s. The remaining component is identified by moving s to the sink and forgetting it.

The first homology of all these spaces is by induction generated by basic classes. Therefore, it remains to show that the classes projecting to the kernel of $H_0(U_1 \cap U_2) \to H_0(U_1) \oplus H_0(U_2)$ are also generated by basic classes.

In $U_1 \cap U_2$ the particle s is trapped on the edge e between z_0 and the central vertex. We can represent each connected component by a configuration where all particles

sit on e. The remaining particles are then distributed to both sides of s. Restricted to the connected components where there is a particle of T on the z_0-side of s, the map

$$U_1 \cap U_2 \hookrightarrow U_2$$

is a homeomorphism onto the corresponding connected components of U_2 because those particles in T prevent s from moving to the sink z_0. The image of that restricted inclusion is disjoint from the image of the remaining components, so to find elements in the kernel of

$$H_0(U_1 \cap U_2) \to H_0(U_1) \oplus H_0(U_2)$$

we can restrict ourselves to the union X of components of $U_1 \cap U_2$ where no element of T is on the z_0-side of s.

The inclusions $X \to U_1$ and $X \to U_2$ map all these connected components to the same component of U_1 and U_2, respectively, because we can use either the sink or the essential vertex to reorder the particles. Therefore, the kernel of the map to $H_0(U_1) \oplus H_0(U_2)$ is generated by differences of distinct ways of putting particles in $S - T$ to the two sides of s, and the lifting process turns these differences into H-classes involving z_0 and the central vertex, proving the claim. □

3.1.1 A basis for configurations in graphs with one essential vertex

The key to proving the induction step is choosing for any fixed leaf e a particular system of bases for the first homology groups $H_1(\mathrm{Conf}_n(G, Z))$ for all n with the following property: if a representative of a basis element has fixed particles on the leaf e then changing the order of these particles should give another basis element, and all these basis elements should be distinct. Furthermore, adding and forgetting fixed particles of representatives of basis elements should again give elements in the chosen system of bases. For the description of such a system of bases, fix the graph G, the set of sinks Z and the leaf e.

We will choose a system of spanning trees T_S (indexed by all finite sets S) in the combinatorial model of $\mathrm{Conf}_S(G, Z)$. As described in Proposition 1.20, this model is a one-dimensional cube complex, i.e. a graph. For each 1-cube ξ in the combinatorial model, the system $T_\bullet$ will have the following properties:

- The edge ξ determines a set F_ξ of fixed particles on the leaf e. The symmetric group $\Sigma_{F_\xi} \leqslant \Sigma_n$ acts on the combinatorial model by precomposition, and we want that the orbit $\Sigma_{F_\xi} \cdot \xi$ is completely contained in either T_S or $G - T_S$.

- Given $s \notin S$ we have a map $\mathrm{Conf}_S(G, Z) \to \mathrm{Conf}_{S \sqcup \{s\}}(G, Z)$ by adding the particle s to the end of the leaf e. Then ξ should be in T_S if and only if the image of ξ under that map is contained in $T_{S \sqcup \{s\}}$.

We now inductively choose the system of spanning trees T_S. For $S = \varnothing$, we define $T_\varnothing = \varnothing$. Given a non-empty set S, complete the forest

$$\bigsqcup_{s \in S} \iota_{e,s}\left(T_{S-\{s\}}\right)$$

to a spanning tree T_S in an arbitrary way (recall that $\iota_{e,s}$ slides in the particle s from the univalent vertex of e, see Definition 1.16). If $S' \subset S$ then $T_{S'}$ appears as subtrees of T_S by adding the particles $S - S'$ to the leaf e in all different orders. While completing this forest we only add edges that have no fixed particles on e, otherwise, one of the trees $T_{S-\{s\}}$ was not maximal in $\mathrm{Conf}_{S-\{s\}}(G, Z)$. This yields a spanning tree T_S of $\mathrm{Conf}_S(G, Z)$, inductively describing spanning trees for all finite sets S with the properties listed above.

This defines a system of bases $\mathcal{B}_\bullet$ of $H_1(\mathrm{Conf}_\bullet(G, Z))$ with the following properties:

- for $\sigma \in \mathcal{B}_S$ the class σ^η given by adding a set of particles T in some order η to the end of the leaf e is an element of $\mathcal{B}_{S \sqcup T}$,
- for $\sigma \in \mathcal{B}_S$ the classes σ^η and $\sigma^{\eta'}$ for two orderings $\eta \neq \eta'$ of T are distinct,
- every $\sigma \in \mathcal{B}_S$ has precisely one *minimal representative* $\sigma_{\min} \in \mathcal{B}_{S'}$ for some $S' \subset S$ such that $(\sigma_{\min})^\eta = \sigma$ for some ordering η of $S - S'$ (meaning that the set S' is minimal with respect to this property) and
- we always have $(\sigma^\eta)_{\min} = \sigma_{\min}$.

Given $\sigma \in \mathcal{B}_S$ and the corresponding minimal cycle C, define S' to be the set of fixed particles of C which are on e. Then $\pi_{S-S'}(\sigma)$ defines the minimal representative $\sigma_{\min} \in \mathcal{B}_{S-S'}$. With this definition it is straightforward to check the four properties described above.

3.1.2 The spectral sequence for the induction step

Let (G, Z) be a tree with loops with any subset of the set of vertices *of valence one* turned into sinks, and let v be an essential vertex which is connected to precisely one other essential vertex w via an edge e. Define the following two open subspaces of G:

$$L := \{x \in G \mid d_G(x, v) < 1\}$$

and

$$K := \{x \in G \mid d_G(x, G - L) < 1\},$$

where d_G is the path metric giving every internal edge of G length 1 and every leaf length $1/2$. In other words, K is the connected component of $G - \{v\}$ containing w, see Figure 3.3.

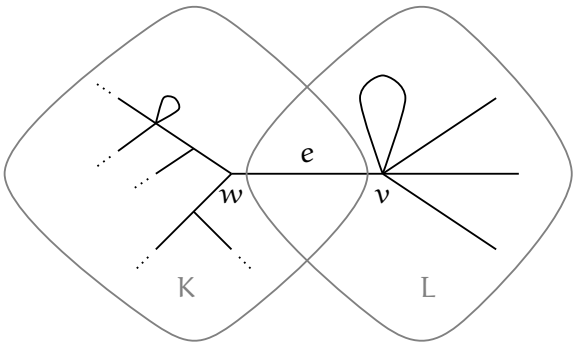

Figure 3.3: The two subgraphs K, L of G.

The intersection $L \cap K$ is the interior of the edge e. The graph K has strictly less essential vertices than G, so by induction we can assume that its configuration spaces (with sinks) of any number of particles are torsion-free and generated by products of basic classes.

As described in Proposition 2.1, construct the open cover of $\mathrm{Conf}_n(G, Z)$ associated to the cover $\{K, L\}$ and look at the corresponding Mayer-Vietoris spectral sequence $E^*_{\bullet,\bullet}$. The open cover has one open set for each map $\phi\colon \underline{\mathbf{n}} \to \{K, L\}$, restricting particle i to the open set $\phi(i)$.

By Proposition 2.5, we have

$$\begin{aligned} U_{\phi_0 \cdots \phi_p} &= U_{\phi_0} \cap \cdots \cap U_{\phi_p} \\ &= \bigcap_{i \in \underline{\mathbf{n}}} \bigcap_{0 \leqslant j \leqslant p} \pi_i^{-1}\left(V_{\phi_j(i)}\right) \\ &\simeq \coprod_{j \in J} \mathrm{Conf}_{S_K^j}(K, Z_K) \times \mathrm{Conf}_{S_L^j}(L, Z_L), \end{aligned}$$

where J is a finite index set, $S_K^j \sqcup S_L^j \subset \underline{\mathbf{n}}$ and Z_K and Z_L are the sinks of K and L, respectively. Remember that in order to get that description, we forgot the particles which are fixed in the interior of the edge between v and w. The order of the particles on this intersection will be important for the face maps.

The E^1-page consists at position (p, q) of the q-th homology of all $(p+1)$-fold intersections of the open sets U_ϕ. By Corollary 2.6, each $E^1_{p,q}$ is given as

$$E^1_{p,q} \cong \bigoplus_{j \in J'} \bigoplus_{q_K + q_L = q} H_{q_K}(\mathrm{Conf}_{S_K^j}(K, Z_K)) \otimes H_{q_L}(\mathrm{Conf}_{S_L^j}(L, Z_L)),$$

where J' is some finite indexing set. Here we used that we know that the configuration spaces of L have free homology. Recall that attached to each of those summands there is an ordering of the particles $\underline{\mathbf{n}} - S_K^j - S_L^j$, which are sitting in the interior of e. The face maps forgetting one of the open sets from a $(p+1)$-fold intersection

yielding a p-fold intersection only affect the particles restricted to the intersection $K \cap L$: for some (but possibly none) of them the restriction is removed, allowing them to move in all of either K or L. Under the identification above, these particles are added to the sets S^j_K or S^j_L and put to the edge e of K or L, respectively, in the order determined by their order on $K \cap L$, see Figure 3.4.

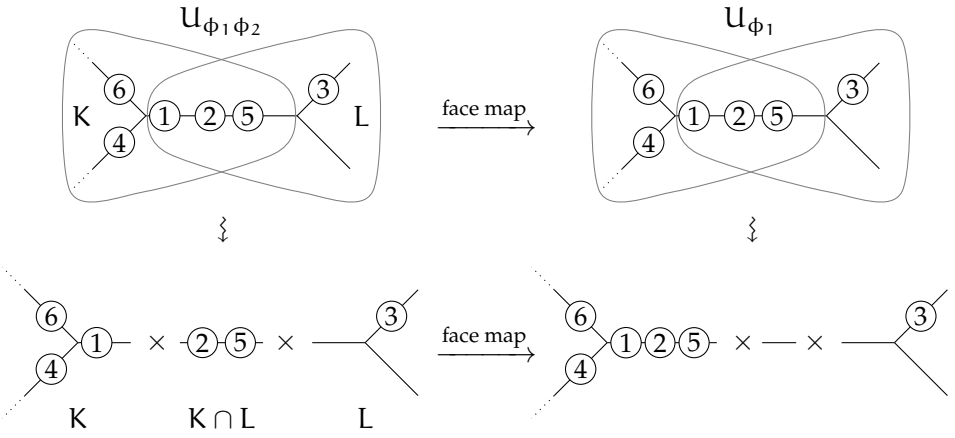

Figure 3.4: The boundary map d_1 of the Mayer-Vietoris spectral sequence under the identification as above. The maps ϕ_1 and ϕ_2 both map $\{1, 4, 6\}$ to K and $\{3\}$ to L. The map ϕ_1 maps the remaining particles $\{2, 5\}$ to K, whereas ϕ_2 maps them to L. The face map to U_{ϕ_1} frees those two particles from $K \cap L$ to K.

Since the configuration space of L is 1-dimensional by Proposition 1.20 these summands of $E^1_{p,q}$ are only non-trivial for $q_L \in \{0, 1\}$. The horizontal boundary map d_1 preserves q_L, so the E^1-page splits into two parts $({}^0E^1, {}^0d_1)$ and $({}^1E^1, {}^1d_1)$ consisting of all direct summands with $q_L = 0$ and $q_L = 1$, respectively. The key points why this helps to compute the infinity page are that (as we will show)

- ${}^1E^2$ is concentrated in the zeroth column,
- we understand ${}^0E^\infty$, and
- the two spectral sequences do not interact.

3.1.3 The homology of ${}^1E^1$

As described in Section 3.1.1, choose a system of bases $\mathcal{B}_\bullet$ for $H_1(\mathrm{Conf}_\bullet(L, Z_L))$ for the edge of L corresponding to e. This determines a direct sum decomposition of the

direct summands of every module ${}^1E^1_{p,q}$ as follows:

$$H_1(\mathrm{Conf}_{S_L^j}(L, Z_L)) \otimes H_{q-1}(\mathrm{Conf}_{S_K^j}(K, Z_K)) \cong \bigoplus_{\sigma \in \mathcal{B}_{S_L^j}} \mathbb{Z}_\sigma \otimes H_{q-1}(\mathrm{Conf}_{S_K^j}(K, Z_K)).$$

Here, $\mathbb{Z}_\sigma$ is the free $\mathbb{Z}$-module on the single generator σ. By the description of the face maps above and the properties of the system of bases, the boundary map ${}^1d_1$ does not change the *minimal* representative of the first tensor factor. Grouping these summands by their corresponding minimal representative σ_0 yields a decomposition of each row ${}^1E^1_{\bullet,q}$ into summands denoted by $(E^1[\sigma_0], d_1^{\sigma_0})$, which is a decomposition *as chain complexes*. We now compute the homology of one of these chain complexes $E^1_{\bullet,q}[\sigma_0]$ for fixed σ_0 and $q \geqslant 0$.

Let a minimal $\sigma_0 \in \mathcal{B}_S$ for some $S \subset \underline{\mathbf{n}}$ be given (i.e. $(\sigma_0)_{\min} = \sigma_0$), then every $\sigma \in \mathcal{B}_{S'}$ appearing in one of the second tensor factors of the modules in the chain complex $E^1_{\bullet,q}[\sigma_0]$ is given by adding fixed particles $S' - S$ to σ_0, putting them in some ordering to the end of e (away from v). Since there are no relations between the different orderings of the particles $S' - S$, we can forget the particles S and replace L by an interval:

Let ${}^K E^*_{\bullet,\bullet}$ be the Mayer-Vietoris spectral sequence for $\mathrm{Conf}_{\underline{\mathbf{n}}-S}(K, Z_K)$ corresponding to the cover $\{K, L\}$ pulled back by the inclusion $K \hookrightarrow G$. The chain complex $E^1_{\bullet,q}[\sigma_0]$ is isomorphic to the chain complex ${}^K E^1_{\bullet,q}$ by forgetting the particles S involved in σ_0 and looking at cycles of the remaining particles, see Figure 3.5.

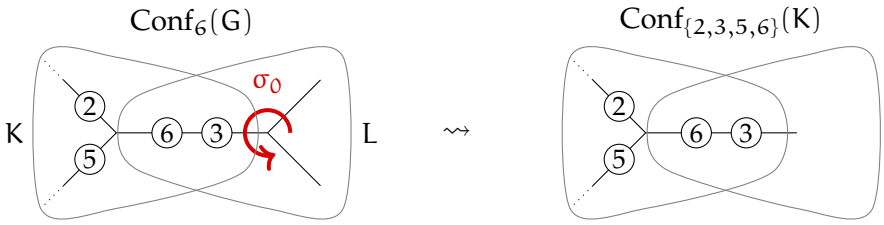

Figure 3.5: Comparing $E^1_{\bullet,q}[\sigma_0]$ to ${}^K E^1_{\bullet,q}$ for $\sigma_0 \in \mathcal{B}_{\{1,4\}}$.

The open cover of K is very special: one of the open sets is the whole space itself. We will now show that because of that, the E^2-page is concentrated in the zeroth column. The open cover of $\mathrm{Conf}_{\underline{\mathbf{n}}-S}(K, Z_K)$ is indexed by maps $\psi \colon \underline{\mathbf{n}} - S \to \{K, L \cap K\}$. For the map ψ_{all} sending everything to K, we have $U_{\psi_{\mathrm{all}}} = \mathrm{Conf}_{\underline{\mathbf{n}}-S}(K, Z_K)$. Hence, for each tuple $(\psi_0, \ldots, \psi_p)$ with $\psi_i \neq \psi_{\mathrm{all}}$ for all i the inclusion

$$U_{\psi_0} \cap \cdots \cap U_{\psi_p} \cap U_{\psi_{\mathrm{all}}} \to U_{\psi_0} \cap \cdots \cap U_{\psi_p}$$

and therefore the face maps

$$H_q(U_{\psi_0} \cap \cdots \cap U_{\psi_p} \cap U_{\psi_{all}}) \to H_q(U_{\psi_0} \cap \cdots \cap U_{\psi_p})$$

are the identity. Notice that precisely one of the $p+2$ face maps with that source lands in an intersection without $U_{\psi_{all}}$. By adding ${}^K d_1$ boundaries we can thus assume that every homology class of the chain complex $({}^K E^1_{p,q}, {}^K d_1)$ has a representative which is trivial in all direct summands $H_q(U_{\psi_0 \cdots \psi_p})$ where none of the ψ_i is ψ_{all}.

The composition of maps

$$\bigoplus_{\substack{\psi_0 < \cdots < \psi_p \\ \exists i : \psi_i = \psi_{all}}} H_q(U_{\psi_0 \cdots \psi_p}) \xrightarrow{{}^K d_1} \bigoplus_{\psi_0 < \cdots < \psi_{p-1}} H_q(U_{\psi_0 \cdots \psi_{p-1}}) \twoheadrightarrow \bigoplus_{\substack{\psi_0 < \cdots < \psi_{p-1} \\ \nexists i : \psi_i = \psi_{all}}} H_q(U_{\psi_0 \cdots \psi_{p-1}}),$$

where the second map sends all direct summands with one of the ψ_i equal to ψ_{all} to zero, is injective by the observation above (actually the images of the direct summands intersect trivially, and restricted to one such summand the map onto its image is given by either the identity or multiplication by -1). In particular, the map ${}^K d_1$ restricted to the intersections including $U_{\psi_{all}}$ is injective (unless we are in the zeroth degree), and the homology is trivial.

Therefore, the homology of $E^1_{\bullet,q}[\sigma_0]$ is zero in degrees $i \neq 0$ and given by

$$\mathbb{Z}_{\sigma_0} \otimes H_{q-1}\left(\mathrm{Conf}_{\underline{n}-S}(K, Z_K)\right)$$

for $i = 0$, which by induction is free and generated by products of basic classes.

In conclusion the homology of ${}^1E^1$ is free, concentrated in the zeroth column and generated by products of basic classes. Denote this bigraded module by $E^\infty[K]$.

3.1.4 The homology of ${}^0E^1$ and the E^∞-page

The other part, ${}^0E^1$, is actually the first page of the Mayer-Vietoris spectral sequence $E^*_{\bullet,\bullet}[G/L]$ of G with $L - e$ collapsed to a sink with respect to the image of the open cover $\{K, L\}^{\underline{n}}$, see Section 2.3. By induction, this spectral sequence $E^*_{\bullet,\bullet}[G/L]$ converges to a free infinity page, and the corresponding homology is generated by products of basic classes.

The E^2-page of our original spectral sequence is hence given by the direct sum of the two bigraded modules $E^2[G/L]$ and $E^\infty[K]$, which differs from $E^2[G/L]$ only in the zeroth column. We will now show that for each $2 \leqslant \ell \leqslant \infty$ the E^ℓ-page is the direct sum of $E^\ell[G/L]$ and $E^\infty[K]$.

For $p > 0$ and $q \geqslant 0$ look at the map d_2 starting in $E^2_{p,q}$. This map is constructed by representing each class in $E^2_{p,q}$ on the chain level (i.e. on the E^0-page), mapping it via the horizontal boundary map to $E^0_{p-1,q}$, lifting it to $E^0_{p-1,q+1}$ and applying the horizontal map again, landing in $E^0_{p-2,q+1}$. The element of $E^2_{p-2,q+1}$ represented

by this cycle is the image of the class we started with under d_2, see Section 2.1.1. The lifting of the cycles in L always connects pairs of distinct orderings of particles on e via a path through the central vertex of L. The end result does not depend on the choice of such a lift, so we always take the following one: choose (once and for all) two edges e_1, e_2 of L that are different from e, then connecting two orderings $\nu \neq \nu'$ of a set $S = \{s_1, \dots, s_m\}$ on e is given by starting with the configuration ν on e, sliding all particles between s_1 and the central vertex to e_2, moving s_1 to e_1, moving the other particles back to e and repeating this for all particles $s_2, \dots, s_m$. Repeating the same for ν' we get two paths which glued together (with opposite orientations) give a path $\gamma[\nu, \nu']$ between the two configurations.

By construction it is clear that $\gamma[\nu, \nu'] + \gamma[\nu', \nu''] = \gamma[\nu, \nu'']$, so the only closed loop arising in such a way is the trivial path. The construction of the image of a class under d_2 as described above produces segments $\gamma[\nu, \nu']$ adding up to a cycle, which hence must be trivial. This shows that d_2 maps to zero in $E^\infty[K]$ and thus that $E^3 \cong E^3[G/L] \oplus E^\infty[K]$. By the same reasoning, this is true for all pages, proving that

$$E^\infty \cong E^\infty[G/L] \oplus E^\infty[K].$$

In conclusion, the E^∞-page is torsion-free and the corresponding homology is generated by products of basic classes.

Proof of Theorem A and Theorem B. For graphs with precisely one vertex of valence at least three and any subset of the set of vertices of valence 1 turned into sinks the theorems follow from Proposition 3.3. By induction on the number of essential vertices, we then use the calculation of the spectral sequence above to prove this for any graph as in the statement of the two theorems with any subset of the set of vertices of valence 1 turned into sinks. In particular, this proves both statements for the case where none of the vertices are sinks. □

Remark 3.7. Notice that we needed to introduce sinks at the leaf vertices because otherwise, we would not be able to describe the infinity page of $E^*_{\bullet,\bullet}[G/L]$, where $L - e$ was collapsed to a point shaped sink. △

Remark 3.8. We will later also use that the embeddings of H can be chosen such that they contain precisely two essential vertices, which can be arranged by splitting an H-graph containing k essential vertices into $k - 1$ of them, each containing exactly two vertices. Also, note that after fixing those two vertices, we can choose the edges of the embedded H-graph arbitrarily: the cycles given by different choices of edges differ by cycles in the stars of the corresponding vertices. △

3.2 The first homology of configurations in general graphs

We know a lot less about the homology of configuration spaces of general graphs. What we *can* describe, however, is a generating system for the *first* homology group

of such configuration spaces.

Theorem C ([CL16, Theorem C, p. 3])**.** *If* G *is any finite graph and* n *a natural number, then the first homology group* $H_1(\mathrm{Conf}_n(G))$ *is generated by basic classes.*

The proof of this statement will be the content of this section. We will formulate the necessary ingredients in a more general way in order to prove the analogous result in the case of graphs with sinks later.

For a connected, finite graph G, we choose distinct edges $e_1, \dots, e_\ell$ such that cutting those edges in the middle yields a tree. Fix identifications of $[0,1]$ with each of the e_i and denote for $x \in [0,1]$ by x_{e_i} the corresponding point on the edge e_i. Then, define the tree K as

$$K = G - \bigcup_{1 \leqslant i \leqslant \ell} [1/3, 2/3]_{e_i},$$

where $[1/3,2/3]_{e_i} = \{x_{e_i} \mid x \in [1/3,2/3]\}$. The idea is now to start with the configuration space of K embedded into the configuration space of G and to release the particles into the bigger graph G one at a time.

For $\Gamma = (G, K)$ recall the definition of $\mathrm{Conf}_{S,T}(\Gamma, Z)$ (Definition 3.6), which restricts the particles $T \subset S$ to the smaller graph K. We will prove that $H_1(\mathrm{Conf}_{S,T}(\Gamma, Z))$ is always generated by basic classes, see Definition 3.1. We will again proceed by constructing an open cover and investigating the Mayer-Vietoris spectral sequence.

Let $\mathrm{Conf}_{S,T}(\Gamma, Z)$ with $S - T$ non-empty be given, then choose an arbitrary element $s \in S - T$ and construct the following open cover: for each i, define two open subsets U_{+e_i} and U_{-e_i} of $\mathrm{Conf}_{S,T}(\Gamma, Z)$ by

$$U_{+e_i} = \left\{ f\colon S \to G \mid f(s) \notin [1/3,2/3]_{e_j} \text{ for } j \neq i \text{ and } f(s) \neq 2/3_{e_i} \right\}$$
$$U_{-e_i} = \left\{ f\colon S \to G \mid f(s) \notin [1/3,2/3]_{e_j} \text{ for } j \neq i \text{ and } f(s) \neq 1/3_{e_i} \right\}.$$

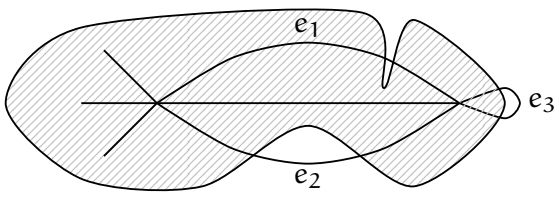

Figure 3.6: The part of G where the particle s is allowed in the open set U_{+e_1}, where e_1 is oriented from left to right.

Let $T' = T \sqcup \{s\}$ and let $\Gamma' = (G - [1/3,2/3]_{e_i}, K)$.

Proposition 3.9 ([CL16, Proposition 4.1, p. 16])**.** *The intersections of those open sets can be identified as follows:*

$$\begin{aligned} U_{\pm e_i} &\simeq \mathrm{Conf}_{S,T'}(\Gamma, Z) \\ U_{-e_i} \cap U_{+e_i} &\simeq \mathrm{Conf}_{S,T'}(\Gamma, Z) \sqcup \mathrm{Conf}_{S-\{s\},T}(\Gamma', Z) \\ U_{\pm e_i} \cap U_{\pm e_j} &\simeq \mathrm{Conf}_{S,T'}(\Gamma, Z). \end{aligned}$$

Any intersection of at least three of those open sets (which are distinct) is homotopy equivalent to $\mathrm{Conf}_{S,T'}(\Gamma, Z)$*.*

The inclusions induced by going from p*-fold intersections to* $(p-1)$*-fold intersections are homotopic to the identity on the components* $\mathrm{Conf}_{S,T'}(\Gamma, Z)$ *and given by adding the particle* s *to* $1/2_{e_i}$ *for the configurations in each component* $\mathrm{Conf}_{S-\{s\},T}(\Gamma', Z)$*. These latter components are not hit by any such inclusion.*

Proof. If the intersection of any number of these open sets contains open sets $U_{\pm e_i}$ and $U_{\pm e_j}$ for $i \neq j$ then the particle s is restricted from entering all $[1/3, 2/3]_{e_i}$, so this intersection is actually precisely the same as $\mathrm{Conf}_{S,T'}(\Gamma, Z)$. Since every intersection of $\geqslant 3$ of those sets contains two such open sets, there are only two cases remaining, namely 1-fold intersections and the intersection $U_{-e_i} \cap U_{+e_i}$.

The space U_{+e_i} is almost the same as $\mathrm{Conf}_{S,T'}(\Gamma, Z)$, the only difference is that the particle s is also allowed in the segment $[1/3, 2/3)_{e_i}$. By sliding s back into the interval $[0, 1/3)_{e_i}$ whenever necessary and moving all particles between 0_{e_i} and s accordingly, we see that this space is homotopy equivalent to $\mathrm{Conf}_{S,T'}(\Gamma, Z)$. The analogous reasoning identifies U_{-e_i}.

The intersection $U_{-e_i} \cap U_{+e_i}$ has two connected components: the component where s is in $(1/3, 2/3)_{e_i}$ and the one where it is in K. The second component is again on the nose equal to $\mathrm{Conf}_{S,T'}(\Gamma, Z)$. Modify the first component by a homotopy moving s to $1/2_{e_i}$ and sliding all other particles on e_i away from s into the intervals $[0, 1/3)_{e_i}$ and $(2/3, 1]_{e_i}$, then forgetting the particle s gives an identification with $\mathrm{Conf}_{S-\{s\},T}(\Gamma', Z)$, proving the first claim.

By our identification above the description of the inclusion maps given by forgetting one of the intersecting open sets is easily deduced. If one of these inclusions would hit a component $\mathrm{Conf}_{S-\{s\},T}(\Gamma', Z)$, then the particle s would need to be on the interval $(1/3, 2/3)_{e_i}$ for some i, which it never is for any triple intersection. □

This allows us to describe generators for the first homology of the configuration space of any finite graph as follows. We formulate this as a separate proposition in order to use it for the case where K is a graph with precisely one essential vertex since this case is needed to prove Theorem B.

Proposition 3.10 ([CL16, Proposition 4.2, p. 17])**.** *Let* G *be a finite graph, let* $K \subset G$ *be a tree defined as above and let* Z *be a subset of the vertex set. If* $H_1(\mathrm{Conf}_S(K, Z))$ *is generated*

by basic classes for all finite sets S then also $H_1(\mathrm{Conf}_{S,T}(\Gamma, Z))$ is generated by basic classes for all pairs of finite sets $T \subset S$, where $\Gamma = (G, K)$.

Proof. We prove this by looking at the spectral sequence constructed from the open cover described above. To prove the statement we only need to show that moving one element out of T preserves the property that the homology is generated by basic classes. We can assume that the configuration space of K is connected since the only case where this is not true is if G is S^1 without sinks, and this case is true by definition. We will now argue by induction on the number of elements in $S - T$. The induction start $S = T$ is precisely that $H_1(\mathrm{Conf}_S(K, Z))$ is generated by basic classes, so we only need to check the induction step.

In the induction step 1-classes arise at $E^\infty_{0,1}$ and $E^\infty_{1,0}$. The module $E^\infty_{0,1}$ is a quotient of $E^1_{0,1}$, which is generated by 1-classes of $U_{\pm e_i} \simeq \mathrm{Conf}_{S,T'}(\Gamma, Z)$, so by induction by classes of the required form.

The chain complex $E^1_{\bullet,0}$ is given by the chain complex of the nerve of the cover (which is a simplex) and one additional copy of $\mathbb{Z}$ for each intersection $U_{+e_i} \cap U_{-e_i}$. Restricted to $H_0(U_{-e_i} \cap U_{+e_i}) \cong \mathbb{Z} \oplus \mathbb{Z}$ the face maps

$$\mathbb{Z} \oplus \mathbb{Z} \cong H_0(U_{-e_i} \cap U_{+e_i}) \to H_0(U_{\pm e_i}) \cong \mathbb{Z}$$

are given by $(x, y) \mapsto \pm(x + y)$. Therefore, all elements $(x, -x)$ are in the kernel of d_1. These elements correspond to S^1 movements of s along the edge e_i: by mapping $U_{-e_i} \cap U_{+e_i} \hookrightarrow U_{-e_i}$ the particle s is allowed to leave $(1/3, 2/3)_{e_i}$ via one of the sides, connecting it to a configuration where s is on the tree K. The other inclusion allows s to leave via the other side, connecting it to that same configuration with s on K. Mapping this to $\mathrm{Conf}_{S,T}(\Gamma, Z)$ yields a cycle where s moves along K and e_i. We can choose a representative such that all other particles are fixed and that this movement follows an embedded circle in G, see Figure 3.7.

Subtracting such kernel elements, we can modify every cycle of $(E^1_{\bullet,0}, d_1)$ such that it is zero in all copies of $H_0(\mathrm{Conf}_{S-\{s\},T}(\Gamma', Z))$. Since the remaining part of the chain complex is the chain complex of a simplex, there are no other 1-classes, concluding the argument. □

Proof of Theorem C. By Theorem B, the group $H_1(\mathrm{Conf}_S(K))$ is generated by basic classes for any finite tree K, so the theorem follows from Proposition 3.10. □

3.3 Homology groups not generated by product cycles

In this section we describe an example of a homology class of the configuration space of a graph that cannot be written as a sum of product classes. This example is taken from [CL16].

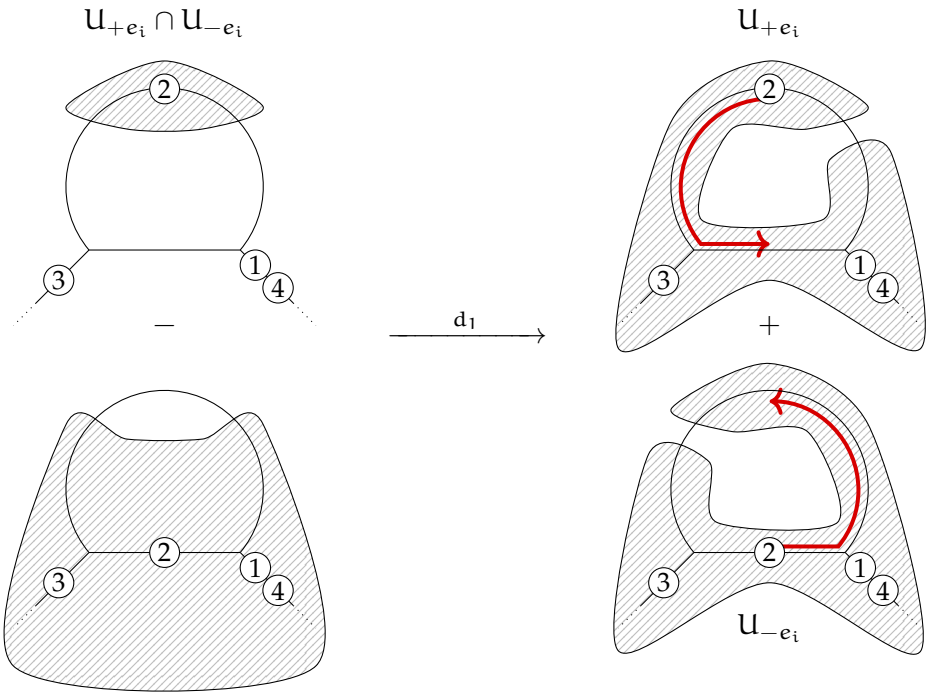

Figure 3.7: The boundary map $H_0(U_{+e_i} \cap U_{-e_i}) \to H_0(U_{+e_i}) \oplus H_0(U_{-e_i})$ producing a circle movement of $s = 2$ along e_i and the tree K.

The easiest example we were able to find so far is a 2-class of $\mathrm{Conf}_3(B_4)$, where B_4 is the banana graph of rank three, i.e. two vertices v, w connected via four edges, see Figure 3.8.

To construct the class, we first construct classes in $\mathrm{Conf}_2(\mathrm{Star}_4)$. Let $S \subset \underline{\mathbf{3}}$ be a set of two particles, then the first homology group of $\mathrm{Conf}_S(\mathrm{Star}_3)$ is one-dimensional, a generator can be represented by a sum of twelve edges, see Figure 1.5.

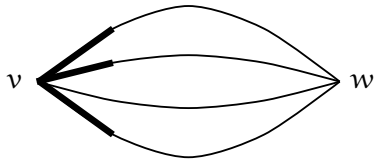

Figure 3.8: Including Star_3 into the banana graph B_4 at v in one of four ways.

Now choose a bijection of $\underline{\mathbf{3}}$ with the leaves of Star_3 and of $\underline{\mathbf{4}}$ with the leaves of Star_4. This defines four 1-cycles in $\mathrm{Conf}_S(\mathrm{Star}_4)$ by including Star_3 into Star_4 in all order-preserving ways (with respect to these identifications). Now we add those four cycles together with the following signs: each inclusion of Star_3 is determined by the edge $i \in \underline{\mathbf{4}}$ that is missed. The 1-cycle corresponding to this i gets the sign $(-1)^i$. This sum is actually equal to zero:

The 1-cells of these cycles are given by one particle moving from one edge to the central vertex and the other particle sitting on another edge. Each such cell appears precisely twice, once for each way of choosing a third edge from the remaining two leaves. If these two remaining leaves are cyclically consecutive in $\underline{\mathbf{4}}$, the corresponding cycles have different signs, otherwise, these two cells inside the 1-cycles appear with different signs, so in both cases, they add up to zero.

Including Star_4 into B_4 (mapping the central vertex to v) gives a sum of four 1-cycles coming from embedding Star_3 into B_4 in different ways (see Figure 3.8). This sum is equal to zero.

Now let t be the third particle, i.e. $S \sqcup \{t\} = \underline{\mathbf{3}}$, then take for each of those four 1-cycles in $\mathrm{Conf}_S(B_4)$ the product of the cycle with the 1-cell moving particle t from the remaining one of the four edges to the vertex w.

Doing this construction for all three choices of S gives a sum of 144 2-cells, and the claim is that this is, in fact, a 2-cycle in the combinatorial model of the configuration space. We can think of this cycle as 12 cylinders, where each cylinder is a 1-cycle in the star of v multiplied with the movement of another particle to the other vertex w, see Figure 3.9. The boundary 1-cells of the cylinders get identified in a certain way.

Let $t \in \underline{\mathbf{3}}$, then one part of the boundary of four of those cylinders is given by the 1-cycles of the particles $\underline{\mathbf{3}} - \{t\}$ with t sitting on w. By construction, those four 1-cycles add up to zero.

It remains to investigate the parts where the particle t moving towards w is in the middle of the edge. These 1-cells are precisely given by two particles sitting in the middle of two edges and a third particle moving from another edge to v. Each such cell appears precisely twice: once for every choice of which one of the fixed particles moves to w and which one belongs to the star movement. By analogous reasoning, these two occurrences have opposite signs, so the total contribution is zero.

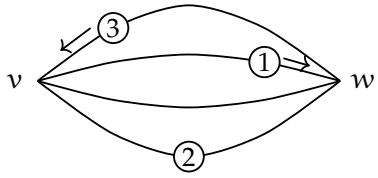

Figure 3.9: Each of the twelve cylinders making up the cycle is given by twelve two cells of this form, where all particles are on different edges.

Thus, the boundary cells of the twelve cylinders add up to zero, yielding a non-trivial cycle. By Proposition 1.20, the combinatorial model is 2-dimensional, so this does not represent the zero class. Notice that there do not exist any product classes since every S^1 generator uses both vertices and there are too few particles for two H-classes or star classes. By looking at the identifications and calculating the Euler characteristic, one sees that the resulting cycle is, in fact, a closed surface of genus 13 embedded into the combinatorial model of the configuration space. In fact, by pushing in 2-cells where strictly less than three edges are involved (starting with those involving only one edge, followed by those involving precisely two edges) and then pushing in the 1-dimensional intervals where particles move to an occupied edge, it is straightforward to show the following:

Proposition 3.11. $\mathrm{Conf}_3(B_4)$ *is homotopy equivalent (equivariantly with respect to the action of the symmetric group* Σ_3*) to a closed orientable surface of genus 13.*

See Figure 3.10 for examples of cells we push in. This is a somewhat degenerate situation, in general we do not expect these configuration spaces to have the homotopy type of a manifold.

This example can now be generalized to give examples of graphs G such that $H_q(\mathrm{Conf}_n(G))$ is not generated by products. We will do this by enlarging the banana graph above and adding particles.

By adding k disjoint S^1 graphs, connecting each of them to v via a single edge and adding k particles we can take the product of the non-product 2-cycle as described above with the k-cycle given by the product of the k particles moving inside the S^1's. This gives a class in the $(k+2)$-nd homology group of the configuration space

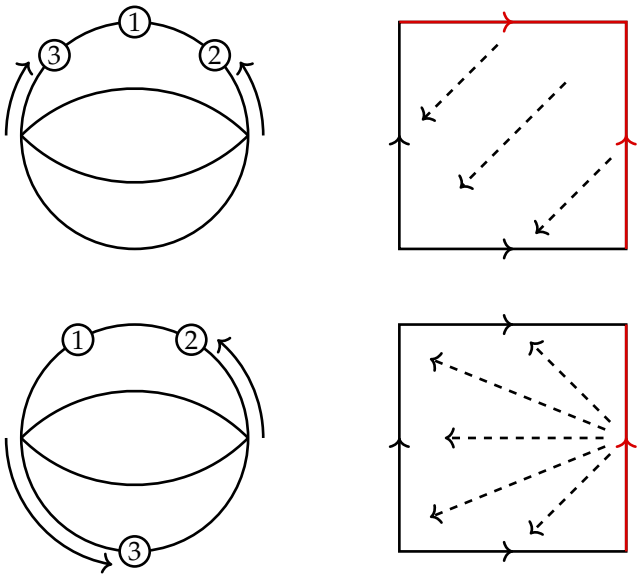

Figure 3.10: Two examples of cells we push in to show that the combinatorial model of $\mathrm{Conf}_3(B_4)$ is homotopy equivalent to the surface of genus 13 we constructed. The horizontal direction in the cube is the movement of particle 2, the vertical direction that of particle 3. The red edges are not incident to any other 2-cell, the dashed arrows indicate how we deform the 2-cell onto the black 1-cells.

of $k+3$ particles in this graph, which by analogous reasoning cannot be written as a sum of product classes. This shows that this phenomenon appears in every homology degree (except for the zeroth and first, of course).

3.4 Graphs with sinks

The description of the homology of configuration spaces of graphs from the previous sections generalizes to the case where an arbitrary subset of the vertex set is turned into sinks. For completeness, we give the missing ingredients to extend our results to this more general setting. Additionally, some of the results used in the proof will be useful in the next chapter.

Theorem D ([CL16, Theorem D, p. 3])**.** *Let* G *be a finite graph and let* Z *be any subset of the vertex set. Then the first homology of* $\mathrm{Conf}_n(G,Z)$ *is generated by basic classes. If* G *is a tree with loops, then* $H_*(\mathrm{Conf}_n(G,Z))$ *is free and generated by products of basic classes.*

To prove this result, we describe the homology of configuration spaces of graphs which are wedged together along a sink.

Proposition 3.12 ([CL16, Proposition 5.1, p. 19])**.** *Let* $(G_1,Z_1),(G_2,Z_2)$ *be based graphs with sinks such that each* Z_i *is a union of vertices and edges of* G_i*, the base points are sinks and* $H_\bullet(\mathrm{Conf}_n(G_i,Z_i))$ *is free for all* $i\in\{1,2\}, n\in\mathbb{N}$*. Then for each* q,n *the homology* $H_q(\mathrm{Conf}_n(G_1\vee G_2, Z_1\cup Z_2))$ *is free and generated by products of homology classes of particles in* G_1 *and* G_2*.*

For the proof we need the following result, choosing a system of bases as in the proof of Theorem A.

Proposition 3.13 ([CL16, Proposition 5.2, p. 20])**.** *Let* $q\in\mathbb{N}$*, let* K *be a based graph and let* Z *be a subset of the vertex set and edges containing the base point* b *such that* $H_q(\mathrm{Conf}_S(K,Z))$ *is free for all finite sets* S*. Then there exists a collection of bases* $\mathcal{B}_S$ *of* $H_q(\mathrm{Conf}_S(K,Z))$ *for each finite set* S *such that*

- *for each* $\sigma\in\mathcal{B}_S$ *and each set* T *the element* $\sigma^T\in H_q(\mathrm{Conf}_{S\sqcup T}(K))$ *given by adding the particles* T *onto the base point is contained in* $\mathcal{B}_{S\sqcup T}$*,*
- *every* $\sigma\in\mathcal{B}_S$ *has precisely one* minimal representative $\sigma_{\min}\in\mathcal{B}_{S'}$ *for* $S'\subset S$ *such that* $(\sigma_{\min})^{S-S'}=\sigma$ *(meaning that the set* S' *is minimal with respect to this property), and*
- *we always have* $\left(\sigma^T\right)_{\min}=\sigma_{\min}$*.*

Proof of Proposition 3.13. For two disjoint sets S and T, define the map

$$\begin{aligned}\iota_T=\iota_T^S\colon \mathrm{Conf}_S(K,Z)&\to\mathrm{Conf}_{S\cup T}(K,Z)\\ \sigma&\mapsto\sigma^T\end{aligned}$$

by adding the particles T onto the base point. Notice that $\pi_S \circ \iota_T^S$ is the identity.

Let $S = \{s_1, \dots, s_\ell\}$ for $\ell \geqslant 1$ be given, then we want to decompose $H_q(\mathrm{Conf}_S(K, Z))$ into direct summands indexed by the set of particles which are not fixed on the base point b. We have that

$$\pi_{\{s_1\}} \circ \iota_{S-\{s_1\}}^{\{s_1\}} \colon H_q(\mathrm{Conf}_{\{s_1\}}(K, Z)) \to H_q(\mathrm{Conf}_S(K, Z)) \to H_q(\mathrm{Conf}_{\{s_1\}}(K, Z))$$

is the identity. Therefore, we have a decomposition into free direct summands

$$\begin{aligned} H_q(\mathrm{Conf}_S(K, Z)) &\cong \iota_{S-\{s_1\}}(H_q(\mathrm{Conf}_{\{s_1\}}(K, Z))) \oplus R_{\{s_1\}} \\ &=: H[s_1]_S \oplus R_{\{s_1\}}. \end{aligned}$$

Now we repeat the same argument for the particle $s_2 \in S$. Since

$$\pi_{\{s_2\}} \colon H_q(\mathrm{Conf}_S(K, Z)) \to H_q(\mathrm{Conf}_{\{s_2\}}(K, Z))$$

maps $H[s_1]_S$ to zero, the composite

$$H_q(\mathrm{Conf}_{\{s_2\}}(K, Z)) \to H_q(\mathrm{Conf}_S(K, Z))/H[s_1]_S \cong R_{\{s_1\}} \to H_q(\mathrm{Conf}_{\{s_2\}}(K, Z))$$

is the identity again, so by the same reasoning, $R_{\{s_1\}}$ has $H[s_2]_S$ as direct summand. Repeating this for all $s_i \in S$ we get a decomposition into free modules

$$\begin{aligned} H_q(\mathrm{Conf}_S(K, Z)) &\cong H[s_1]_S \oplus \cdots \oplus H[s_\ell]_S \oplus R[2]_S \\ &=: H[1]_S \oplus R[2]_S. \end{aligned}$$

Notice that all classes with a representative having $n - 1$ fixed particles on b map to zero in the $R[2]_S$-summand.

Now take the subset $\{s_1, s_2\} \subset S$ and look at the composite map

$$\pi_{\{s_1,s_2\}} \circ \iota_{S-\{s_1,s_2\}} \colon H_q(\mathrm{Conf}_{\{s_1,s_2\}}(K, Z)) \to H_q(\mathrm{Conf}_{\{s_1,s_2\}}(K, Z)),$$

which is again the identity. By taking quotients by $H[1]_{\{s_1,s_2\}}$ and $H[1]_S$, this gives the following composite, which is still the identity:

$$R[2]_{\{s_1,s_2\}} \to R[2]_S \to R[2]_{\{s_1,s_2\}}.$$

Therefore, we get a direct sum decomposition

$$\begin{aligned} R[2]_S &\cong \iota_{S-\{s_1,s_2\}}(R[2]_{\{s_1,s_2\}}) \oplus R_{\{s_1,s_2\}} \\ &=: H[s_1, s_2] \oplus R_{\{s_1,s_2\}}. \end{aligned}$$

Repeating the arguments from above with all 2-element subsets of S gives

$$R[2]_S \cong H[s_1, s_2]_S \oplus H[s_1, s_3] \oplus \cdots \oplus H[s_{\ell-1}, s_\ell] \oplus R[3]_S$$

and therefore a decomposition into free modules as follows:

$$H_q(\mathrm{Conf}_S(K,Z)) \cong H[1]_S \oplus H[2]_S \oplus R[3]_S.$$

Continuing this process we eventually get a decomposition

$$H_q(\mathrm{Conf}_S(K,Z)) \cong H[1]_S \oplus \cdots \oplus H[\ell-1]_S \oplus H[\ell]_S,$$

where $H[\ell]_S = H[S]_S := R[\ell]_S$.

Choosing arbitrary bases for $H[T]_T$ for all finite sets T determines bases of $H[T]_S$ for all $T \subset S$ and therefore a collection of bases $\mathcal{B}_\bullet$ for $H_q(\mathrm{Conf}_\bullet(K,Z))$. By construction, this collection has the required three properties. □

Proof of Proposition 3.12. This is very similar to the proof of Theorem A and Theorem B, so we will omit a few details. Cover the graph $G := G_1 \vee G_2$ by two open sets U_1, U_2, where U_i is defined to be the union of G_i and the open ball of radius $1/2$ around the base point of G for the path metric where each edge has length 1. The intersection $U_1 \cap U_2$ is a star graph Star_b whose central vertex is a sink, so its configuration space is contractible. We now look at the Mayer-Vietoris spectral sequence corresponding to the cover of the configuration space, see Proposition 2.1. The entries on the E^1-page are given by direct sums of the q-th homology of $(p+1)$-fold intersections of open sets, and by Proposition 2.5 each module $E^1_{p,q}$ is a direct sum of modules of the form

$$\begin{aligned} &H_q(\mathrm{Conf}_{S_1}(G_1,Z_1) \times \mathrm{Conf}_{S_b}(\mathrm{Star}_b, Z_{\mathrm{Star}_b}) \times \mathrm{Conf}_{S_2}(G_2,Z_2)) \\ &\cong \bigoplus_{q_1+q_2=q} H_{q_1}(\mathrm{Conf}_{S_1}(G_1,Z_1)) \otimes H_{q_2}(\mathrm{Conf}_{S_2}(G_2,Z_2)) \end{aligned}$$

for some $S_1 \sqcup S_b \sqcup S_2 \subset \underline{n}$.

Using Proposition 3.13, choose systems of bases $\mathcal{B}_\bullet^{q,i}$ of $H_q(\mathrm{Conf}_\bullet(G_i,Z_i))$, determining bases for all entries on the E^1-page by the identification above. Then each row of the E^1-page splits into a sum of chain complexes indexed by pairs of minimal representatives in $\mathcal{B}_\bullet^{q,1}$ and $\mathcal{B}_\bullet^{q,2}$: for each such pair of minimal representatives take the graded submodule of the q-th row generated by all pairs in $\mathcal{B}_{S_1}^{q,1} \times \mathcal{B}_{S_2}^{q,2}$ for any $S_1 \sqcup S_2 \subset \underline{n}$ having these minimal representatives. This determines a direct sum decomposition as graded modules. Since the horizontal boundary map under the identification above only adds fixed particles to the base point and therefore does not change the minimal representatives, this, in fact, is a decomposition of the q-th row *as chain complexes*, see also the proof of Theorem A and Theorem B.

Now look at such a chain complex $\mathcal{C}_{\sigma_1\sigma_2}$ for some pair of minimal representatives $\sigma_1 \in \mathcal{B}_{S_1}^{q,1}$ and $\sigma_2 \in \mathcal{B}_{S_2}^{q,2}$. For $\sigma \in \mathcal{B}_T^{q,1}$ with minimal representative σ_1 there exists a representative of σ where the particles $T - S_1$ are fixed on the base point, and the analogous statement holds for σ_2. Hence, the generators in the chain complex $\mathcal{C}_{\sigma_1\sigma_2}$

only differ in the subsets of particles in $\underline{\mathbf{n}} - S_1 - S_2$ sitting on b which they put into U_1, U_2 and the intersection $U_1 \cap U_2$.

Let ${}^bE^*_{\bullet,\bullet}$ be the Mayer-Vietoris spectral sequence for $\mathrm{Conf}_{\underline{\mathbf{n}}-S_1-S_2}(\mathrm{Star}_b, Z_{\mathrm{Star}_b})$ for the open cover $\{U_1, U_2\}$ pulled back via $\mathrm{Star}_b \hookrightarrow G$. Then $\mathcal{C}_{\sigma_1\sigma_2}$ is isomorphic to the chain complex ${}^bE^1_{\bullet,0}$, see Section 3.1.3. Since $\mathrm{Conf}_{\underline{\mathbf{n}}-S_1-S_2}(\mathrm{Star}_b, Z_{\mathrm{Star}_b})$ is contractible and ${}^bE^1_{p,q}$ is trivial for $q > 0$, the homology of this chain complex is $\mathbb{Z}$ in degree zero and 0 else.

The same is thus true for $\mathcal{C}_{\sigma_1\sigma_2}$, whose zeroth column is then generated by the product of σ_1 and σ_2 with the remaining particles sitting on b. □

We are now ready to prove Theorem D.

Proof of Theorem D. The proofs of Theorem A and Theorem B worked for any subset of the set of vertices *of valence one* turned into sinks. Using Proposition 3.12, the claims follow for any subset of the vertex set turned into sinks by wedging together such graphs.

Therefore, the group $H_1(\mathrm{Conf}_S(K, Z))$ is generated by basic classes for any finite tree K, so also $H_1(\mathrm{Conf}_S(G, Z))$ is generated by basic classes by Proposition 3.10. □

Similarly to the non-sink case, the homology of configuration spaces of graphs with sinks is in general not generated by product classes. The type of non-product classes, however, is rather different. We will now construct such a class which does not have a correspondence in the non-sink case. This example is taken from [CL16].

Let B_3 be the theta graph, i.e. two vertices v and w connected by three edges, and let $Z = \{v, w\}$. Then there exists a surface of genus two embedded into the combinatorial model of $\mathrm{Conf}_2(B_3, Z)$, constructed as follows.

Let e_1, e_2 and e_3 be the edges of the theta graph, then look at the torus T_1 in $B_3 \times B_3$ given by particle 1 moving along the S^1 consisting of e_1 and e_2 and particle 2 moving along the S^1 consisting of e_1 and e_3. Let T_2 be the same torus with the particles 1 and 2 exchanged.

The tori T_1 and T_2 are not contained in the configuration space of B_3: for each T_i precisely one 2-cell is outside the configuration space, namely the cell where both 1 and 2 are moving along e_1 at the same time. The difference of T_1 and T_2, however, erases precisely that 2-cell and produces a 2-cycle of genus two in $\mathrm{Conf}_2(B_3, Z)$, see Figure 3.11. Since there are no three-cells, this represents a non-zero homology class. There is no pair of embedded circles that does not share an edge, so in particular there are no product classes, showing that this cycle cannot be written as a sum of products. In fact, this cycle is precisely the sum of all cubes in the combinatorial model, so the combinatorial model of $\mathrm{Conf}_2(B_3, Z)$ is precisely this surface of genus two.

This very different behavior hints at the fact that the notion of sinks might not be as useful for the investigation of general graphs, at least compared to the tree with loops case. In the next chapter, however, we will see how to use sinks to deduce representation stability for the second homology group of configuration spaces of graphs (where we stabilize the graph), even though we are not able to give a concrete description of all homology classes.

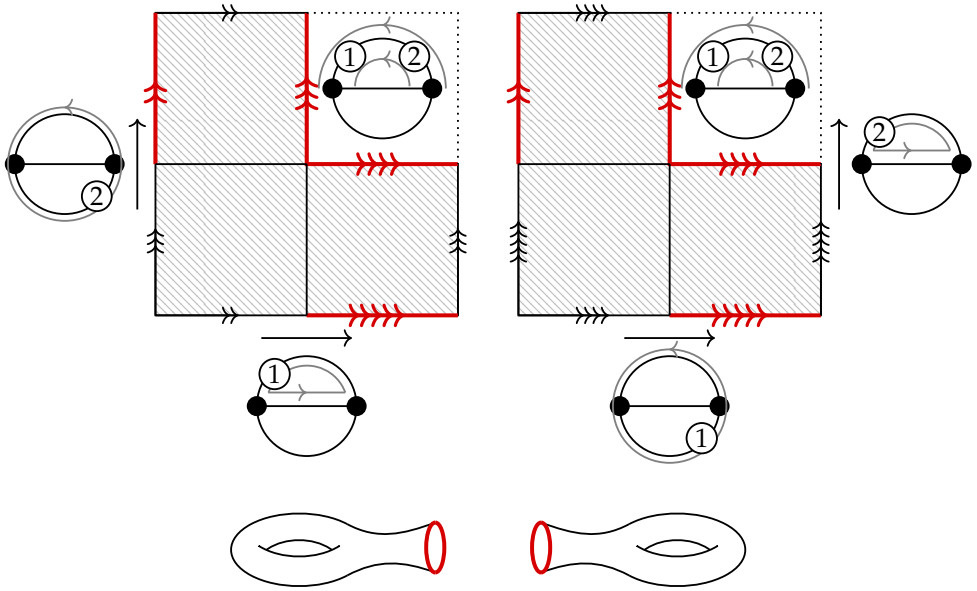

Figure 3.11: The surface of genus two in the theta graph B_3 with two sinks. All black opposite intervals are identified pairwise as indicated, giving two tori with boundary consisting of the red intervals. Gluing the red intervals together as indicated produces a surface of genus 2. The dotted lines indicate a 2-cell which is only there in the product $B_3 \times B_3$, but not in the configuration space. The lower left corner corresponds in both tori to the configuration where both particles are on the left sink.

Chapter 4

Representation stability for configuration spaces of graphs

In this chapter we will prove stability results for configuration spaces of graphs. We briefly review the concept of *representation stability*. Then, we prove representation stability for a fixed number of particles in a growing graph, where the symmetric groups act on the graph. In the final part of this chapter we keep the graph fixed and prove that if this graph has high vertex connectivity the map forgetting the last particle induces a representation stable sequence of first cohomology groups $H^1(\mathrm{Conf}_\bullet(G);\mathbb{Q})$.

4.1 Representation stability and $\mathrm{FI}^{\times \ell}$-modules

In [CF13], Church and Farb introduced the concept of *representation stability*. We now recall the concept in the case of the symmetric groups Σ_k, for more details see [CF13, Section 2.3, p. 19].

Let $\{V_k\}_{k\in\mathbb{N}}$ be a sequence of representations, where V_k is a Σ_k-representations over $\mathbb{Q}$, with linear maps

$$\phi_k \colon V_k \to V_{k+1}$$

which are homomorphisms of $\mathbb{Q}\Sigma_k$-modules. Here we consider V_{k+1} as $\mathbb{Q}\Sigma_k$ module by the standard inclusion $\Sigma_k \hookrightarrow \Sigma_{k+1}$.

To describe stability for such a sequence, we need to compare Σ_k-representations to $\Sigma_{k'}$-representations for $k' > k$. Recall that the irreducible representations of Σ_k over the rational numbers are in one to one correspondence to partitions λ of k. Given a partition $\lambda = (\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_\ell > 0)$ of k — which means $\lambda_1 + \cdots + \lambda_\ell = k$ — define for $k' - k \geqslant \lambda_1$ the irreducible $\Sigma_{k'}$-representation $V(\lambda)_{k'}$ to be the one corresponding to the partition $(k'-k, \lambda_1, \ldots, \lambda_\ell)$. Each irreducible representation of $\Sigma_{k'}$ can be written like this for a unique partition λ. If one thinks of the partition λ as

a Young diagram with λ_i boxes in the i-th row, then two irreducible representations of Σ_k and $\Sigma_{k'}$ for $k < k'$ are viewed as "the same" irreducible representation if the Young diagram of the second one can be obtained by adding boxes to the first row of the Young diagram of the first representation. For more details, see [CF13, Section 2.1, p. 14] and [FH91].

Definition 4.1 ([CF13, Definition 2.3, p.20])**.** The sequence $\{V_k\}$ is *(uniformly) representation stable* if, for sufficiently large k, each of the following conditions holds.

- $\phi_k \colon V_k \to V_{k+1}$ is injective.
- The $\mathbb{Q}\Sigma_{k+1}$ submodule generated by $\phi_k(V_k)$ is equal to V_{k+1}.
- Decompose each V_k into irreducible representations

$$V_k = \bigoplus_{\lambda} V(\lambda)_k^{\oplus c_{\lambda,k}}$$

  with multiplicities $0 \leqslant c_{\lambda,k} \leqslant \infty$. Then there exists an $N \geqslant 0$ such that for each λ, the multiplicity $c_{\lambda,k}$ is independent of $k \geqslant N$. △

This reduces the description of the infinite sequence of Σ_k-representations to a finite calculation.

In [CEF15], Church-Ellenberg-Farb introduced the notion of **FI**-modules, which we now recall. Let **FI** be the category with objects all finite sets and morphisms all injective maps. We often consider the skeleton of this category given by the restriction to the finite sets $\underline{\mathbf{n}} := \{1, \dots, n\}$ for $n \geqslant 0$.

Definition 4.2. Let R be a commutative ring. An $R[\mathbf{FI}]$-module $V_\bullet$ is a functor

$$V_\bullet \colon \mathbf{FI} \to \mathrm{RMod}.$$

It is said to be *finitely generated in degree* ℓ if there exists a finite set X of elements in

$$\bigsqcup_{\substack{S \in \mathbf{FI} \\ |S| \leqslant \ell}} V_S,$$

such that the smallest sub-**FI**-module containing all these elements is $V_\bullet$. Here, $|S|$ is the cardinality of S. △

For finite-dimensional representations over fields of characteristic 0, finitely generated **FI**-modules produce representation stable sequences in the following way:

Theorem 4.3 ([CEF15, Theorem 1.13, p. 8])**.** *An* **FI***-module* $V_\bullet$ *over a field of characteristic 0 is finitely generated if and only if the sequence* $k \mapsto V_{\underline{k}}$ *is representation stable and each* $V_{\underline{k}}$ *is finite-dimensional.*

This result reduces the uniform decomposition of the representations $V_{\underline{k}}$ to finding a finite set of generators. Furthermore, Church-Ellenberg-Farb proved that the dimension of representation stable sequences grows polynomially:

Theorem 4.4 ([CEF15, Theorem 1.5, p. 4])**.** *Let $V_\bullet$ be an* **FI***-module over a field of characteristic 0. If $V_\bullet$ is finitely generated then the sequence of characters $\chi_{V_\bullet}$ is eventually polynomial. In particular,* $\dim V_{\underline{k}}$ *is eventually polynomial in* k*.*

In order to describe stabilization in multiple "directions" we look at the product category $\mathbf{FI}^{\times \ell}$ consisting of $\ell \geqslant 1$ copies of the category **FI**. An $\mathbf{FI}^{\times \ell}$-module is then a functor $\mathbf{FI}^{\times \ell} \to \mathrm{RMod}$, the notion of finite generation is defined analogously.

To define such a module, it is sufficient to define it on the skeleton consisting of the objects $(\underline{\mathbf{j_1}}, \dots, \underline{\mathbf{j}_\ell})$ for $j_i \in \mathbb{N}$ and the morphisms between them.

Clearly, if V is a finitely generated $\mathbf{FI}^{\times \ell}$-module and $F\colon \mathbf{FI} \to \mathbf{FI}^{\times \ell}$ is any functor, then the **FI**-module $F^*V := V \circ F$ is finitely generated: each component of F is either eventually constant or unbounded. The diagonal is one example of such a functor.

4.2 Stability for trees with loops and the first homology group

In the next two sections we investigate the stabilization behavior of configuration spaces of graphs when we keep the number of particles fixed and stabilize the graph. We will first define precisely what we mean by graph stabilization and then show that by the description of the generators for the homology of configuration spaces of trees with loops in Theorem B we can show that these homology groups indeed stabilize.

Let G_0 be a finite graph and let $K_i \subset G_i$ for $1 \leqslant i \leqslant \ell$ be pairs of finite graphs such that each K_i is also a subgraph of G_0. Denote by $\underline{\Gamma} = \underline{\Gamma}_{G_0} := \{(K_1, G_1), \dots, (K_\ell, G_\ell)\}$. Let $\mathbf{G} = \mathbf{G}_{\underline{\Gamma}}\colon \mathbf{FI}^{\times \ell} \to \mathbf{Top}$ be given by

$$\mathbf{G}_{\underline{\Gamma}}(\mathbf{j_1}, \dots, \mathbf{j}_\ell) := G_0 \sqcup_{K_1} G_1^{\sqcup j_1} \cdots \sqcup_{K_\ell} G_\ell^{\sqcup j_\ell},$$

i.e. by gluing the copies of the graphs G_i to G_0 via the shared subgraph K_i. To define the images of morphisms, notice that each summand G_i can be labeled by a number between 1 and j_i. For a map $\phi\colon \underline{\mathbf{j_i}} \hookrightarrow \underline{\mathbf{j}_i'}$ we define the induced map to send the summand with label $m \in \underline{\mathbf{j_i}}$ to the summand with label $\phi(m)$ via the identity.

Meta Question. *For which* G_i, K_i, q *and abelian group* A *is the* $\mathbf{FI}^{\times \ell}$*-module*

$$\mathbf{H}^A_{q,n}\underline{\Gamma} := H_q(\mathrm{Conf}_n(\mathbf{G}_{\underline{\Gamma}}); A)$$

finitely generated?

For $A = \mathbb{Z}$ we also write $\mathbf{H}^{\underline{\Gamma}}_{q,n}$. Part of this question asks how "local" the homology of configuration spaces in graphs is. In the case of trees with loops we proved a

rather strong kind of locality in Theorem B by describing an explicit generating system of products of 1-classes, and by refining the generating system we can prove in that case that the answer to the question above includes all trees with loops.

Theorem E. *If each of the graphs G_i for $0 \leqslant i \leqslant \ell$ is a tree with loops, then $\mathbf{H}^{\Gamma}_{q,n}$ is finitely generated in degree $(\zeta, \zeta, \dots, \zeta)$ for each $q, n \in \mathbb{N}$, where $\zeta = \zeta_{n,q} = \min\{2n, n + 3q\}$.*

Corollary 4.5. *In the same situation as in the theorem above choose any functor*

$$F\colon \mathbf{FI} \to \mathbf{FI}^{\times \ell},$$

then the **FI***-module $\mathbf{H}^{\Gamma}_{q,n} \circ F$ is finitely generated. In particular, the sequence $\mathbf{H}^{\mathbb{Q}}_{q,n}\underline{\Gamma} \circ F$ is representation stable and therefore the dimension of the sequence of vector spaces is eventually polynomial.*

Corollary 4.6. *Let G, K be finite trees with loops with base point and define*

$$G_k := G \vee \underbrace{\underline{K} \vee \cdots \vee \underline{K}}_{k \text{ times}}.$$

Then the **FI***-module $H_q(\mathrm{Conf}_n(G_\bullet))$ is finitely generated. In particular, the sequence $H_q(\mathrm{Conf}_n(G_\bullet); \mathbb{Q})$ is representation stable and therefore the dimension of the sequence of vector spaces is eventually polynomial in k.*

Remark 4.7. The fact that the dimension of the sequence in the previous corollary is *bounded from above* by a polynomial can be seen more easily: The number of edges of G_k and the maximal valence of the vertices in the graph G_k are both polynomial in k. By the description of the combinatorial model in [Lü14], the number of cells is (for fixed number of particles) bounded from above by a polynomial in these two numbers. Therefore, the number of ℓ-cells in the combinatorial model and thus the maximal dimension of the ℓ-th homology group is at most polynomial in k.

The result, however, shows more: it shows that the dimension of the homology is eventually *equal* to a polynomial in k, and additionally that the representations of the symmetric group eventually stabilize. See Section 4.1 for more details. △

For a general graph we currently only have a description of generators for the *first* homology group of its configuration space, see Theorem C. This description allows us to answer the question completely for $q = 1$:

Theorem F. *For any choice of graphs G_i and K_i the* **FI***-module $\mathbf{H}^{\Gamma}_{1,n}$ is finitely generated in degree $(n + 3, n + 3, \dots, n + 3)$ for each $n \in \mathbb{N}$.*

In order to prove those two theorems we first prove Corollary 4.6 for star graphs by hand.

Proposition 4.8. *Corollary 4.6 is true for* G *the point and* K *the interval* $[0,1]$ *with 0 as base point, i.e. the sequence of star graphs with increasing number of leaves. In fact, the homology is generated by cycles meeting at most* $n+3$ *many copies of* K*, so the* **FI***-module is generated in degree* $n+3$*.*

Remark 4.9. For $n = 2$ the argument presented below is easily modified to show that $H_1(\mathrm{Conf}_2(G_\bullet))$ is generated in degree $n+2 = 4$. Since $n+3 \leqslant 2n$ for $n > 2$, this shows that $H_1(\mathrm{Conf}_n(G_\bullet))$ is generated in degree $2n$, which will be used in the proof of Theorem E. △

Proof of Proposition 4.8. The combinatorial model of this configuration space is a graph, so we only need to consider 1-cycles. Choose any subgraph $\mathrm{Star}_3 \subset \mathrm{Star}_k$. Let C be a 1-cycle, then the claim is that we can write C as a sum of cycles where each particle uses at most one edge outside of Star_3.

Let x be a particle and choose a 0-cube v of C where x sits on the vertex of the star. If this does not exist, then x is fixed and therefore uses at most one edge. Now move along a path γ of 1-cubes in the cycle until x sits on the vertex again and there exists a continuation such that the next 1-cube would move x onto an edge of $\mathrm{Star}_k - \mathrm{Star}_3$ for the *second time*. We denote the corresponding terminal 0-cube of γ by v'. Now choose the following path γ' back to v, during which x always stays in Star_3: move x onto an edge e_1 of Star_3 and keep it there. Follow γ back ignoring the movement of x and using the connectedness of the configuration space of Star_3 to move x out of the way if other particles need to move along e_1. Finally, move x back to the vertex.

This decomposes C into two cycles: the cycle $\gamma\gamma'$ and C with γ replaced by γ'^{-1}. In the first of those two cycles the particle x only visits one edge not in Star_3. Continuing this process, we eventually exhaust all edges of C and get a sum decomposition of C where in each summand x visits only Star_3 and at most one additional edge.

Since we did not increase the number of edges outside of Star_3 visited by any other particle, we can repeat this for every x and get a sum decomposition of C of the required form.

Consequently, for each $N \geqslant n+3$ we can generate $H_1(\mathrm{Conf}_n(G_N))$ by cycles such that each one of them is supported in some subgraph $\mathrm{Star}_{n+3} \hookrightarrow G_N$. Therefore, the $\mathbb{Z}\Sigma_N$-span of the image of the map

$$H_1(\mathrm{Conf}_n(G_{n+3})) \to H_1(\mathrm{Conf}_n(G_N))$$

is the whole module and the **FI**-module $H_1(\mathrm{Conf}_n(G_\bullet);\mathbb{Z})$ is finitely generated in degree $n+3$. □

Proof of Theorem E. Let $n > 1$ and let $(k_1,\dots,k_\ell)$ be such that each k_i is at least $\zeta = \min\{2n, n+3q\}$. By Theorem B, the homology of $\mathrm{Conf}_n(\mathbf{G}(k_1,\dots,k_\ell))$ is generated by products of basic cycles. By Remark 3.8, we can assume that the

embedded H-graphs contain exactly two vertices because $k_i > 3$ and therefore the valence of all internal vertices is at least three. In the following, we will say that a particle meets a copy of some G_i if it moves into the part not contained in G_0, namely the part $G_i - K_i$.

Each **H-class** meets at most 3 copies of each G_i by Remark 3.8. Since each H-class consists of $m \geqslant 2$ particles, we have $3 \leqslant \zeta_{m,1}$.

Each **star class** with m particles can be written as a linear combination of generators such that each is using only $\zeta_{m,1}$ different edges by Proposition 4.8 and Remark 4.9. Therefore, each summand visits at most $\zeta_{m,1}$ distinct copies of each of the G_i.

The only embedded copies of S^1 are given by self loops at one of the vertices, so each **S^1-class** meets at most one of the copies of one of the G_i.

Each of the non-moving particles meets at most one of the copies. Hence, we can generate the whole homology by classes which each meet at most

$$\zeta_{m_1,1} + \cdots + \zeta_{m_q,1} + (n - m_1 - \cdots - m_q) \leqslant \min\{2n, n + 3q\} = \zeta_{n,q}$$

different copies of each of the G_i. This implies that the $\mathbb{Z}[\Sigma_{k_1} \times \cdots \times \Sigma_{k_\ell}]$-span of the image of

$$H_q(\mathrm{Conf}_n(\mathbf{G}(\zeta, \ldots, \zeta))) \to H_q(\mathrm{Conf}_n(\mathbf{G}(k_1, \ldots, k_\ell)))$$

is the whole module, finishing the proof. □

Proof of Theorem F. By Theorem C, the homology group $\mathbf{H}^{\Gamma}_{1,n}$ is generated by basic classes. Fix $(k_1, \ldots, k_\ell)$ such that $k_i \geqslant n + 3$ for each i. Each H-class meets at most three copies of each G_i by Remark 3.8. Star classes involving k particles can be written as sums of other star classes, each meeting at most $k + 3$ copies of each of the G_i by Proposition 4.8. Every S^1-class can be written as a sum of S^1-classes such that each of them meets at most two copies of each of the G_i: choose a spanning tree for each connected component of $\mathbf{G}(1, \ldots, 1) \subset \mathbf{G}(k_1, \ldots, k_\ell)$ and extend it to spanning trees for the connected components of $\mathbf{G}(k_1, \ldots, k_\ell)$. The inclusion $\mathbf{G}(1, \ldots, 1) \hookrightarrow \mathbf{G}(k_1, \ldots, k_\ell)$ is a π_0-isomorphism, so this construction ensures that this forest restricted to the union of $\mathbf{G}(1, \ldots, 1)$ and a copy of one of the G_i still gives a spanning forest. The cycles corresponding to the edges outside of that spanning forest thus stay inside this copy and $\mathbf{G}(1, \ldots, 1)$, so they meet at most two copies of each G_i.

The non-moving particles meet at most one copy each, so each class can be written as a sum of classes meeting at most $n + 3$ copies of each of the G_i. □

4.3 Stability for banana graphs

In this section we will prove stabilization for *banana graphs*, which will be the main input for proving stabilization of the second homology in general, see Theorem H.

Definition 4.10. For $k \in \mathbb{N}$, the graph B_k given by two vertices v and w connected via k edges is called the *banana graph with* k *edges*. This graph has an action of the symmetric group Σ_k given by permuting the edges. △

In fact, we will view the collection of spaces $B_\bullet$ as an **FI**-space in the canonical way.

Theorem G. *For each* $n \in \mathbb{N}$ *and* $q \in \mathbb{N}$ *the* **FI***-module* $H_q(\mathrm{Conf}_n(B_\bullet))$ *is finitely generated in degree* $n+6$*. The* **FI***-module* $H_q(\mathrm{Conf}_n(B_\bullet, Z))$ *is finitely generated in the same degree for any* $Z \subset \{v, w\}$ *and* $q \leqslant 3$*.*

Remark 4.11. It is easier to only see that the dimension of $H_\bullet(\mathrm{Conf}(B_k))$ is eventually polynomial in k. From Theorem C, we see that the first homology is generated by basic classes – by star classes, H-classes and S^1-classes. By Theorem F, the **FI**-module $H_1(\mathrm{Conf}_n(B_\bullet))$ is generated in degree $n+3$, so $\dim H_1(\mathrm{Conf}_n(B_k); \mathbb{Q})$ is eventually polynomial in k. The Euler characteristic of $\mathrm{Conf}_n(B_k)$ is polynomial in k, which can be seen by the formulas for the number of cells of the combinatorial model in [Lü14, Section 3.4, p. 38]. Since the zeroth homology is $\mathbb{Z}$ for $k \geqslant 3$, this shows that the dimension of the second homology group is also eventually polynomial in k. This, however, is weaker than saying that the **FI**-module given by the second homology is finitely generated. △

We will reduce Theorem G to finding a good set of generators for the homology $H_*(\mathrm{Conf}_n(B_k, \{v, w\}))$ of banana graphs with two sinks, and then describe such a set of generators, see Proposition 4.13 and Proposition 4.14.

Remark 4.12. The case $Z = \varnothing$ will also follow from Theorem H, but we still give the explicit proof here to illustrate the techniques on a small example. △

Proof of Theorem G. By Theorem F we only have to prove the statement for $q \geqslant 2$. If $|Z| = 1$, the space in question is homotopy equivalent to a 1-dimensional complex, so there is nothing to prove. The case $|Z| = 2$ will follow from Proposition 4.13 and Proposition 4.14, so it remains to handle the case $Z = \varnothing$ and $q = 2$. Furthermore, we can assume $n \geqslant 2$ because for $n = 1$ the second homology is trivial.

For fixed $k \in \mathbb{N}$, look at the open cover $\mathcal{V} = \{V_v, V_w\}$ with

$$V_v := B_k - \{w\} \qquad \text{and} \qquad V_w := B_k - \{v\}.$$

Both of these open sets are star graphs (with missing vertices of valence 1). Denote by $E^*_{\bullet,\bullet}[Z]$ the spectral sequence defined via the open cover $\mathcal{V}^{\underline{n}}$ of $\mathrm{Conf}_n(B_k, Z)$, as described in Section 2.1. For $Z = \varnothing$, we write $E^*_{\bullet,\bullet} = E^*_{\bullet,\bullet}[\varnothing]$.

We want to prove that the **FI**-module $H_2(\mathrm{Conf}_n(B_\bullet))$ is finitely generated, so we are interested in the groups $E^\infty_{0,2}$, $E^\infty_{1,1}$ and $E^\infty_{2,0}$. The first one, $E^\infty_{0,2}$, is a quotient of $E^1_{0,2}$ and therefore generated by products of star classes. By Proposition 4.8, every

such product of n_1 and n_2 particles forming a star class can be written as sums of such products meeting at most $n_1 + 3 + n_2 + 3$ edges. The remaining $n - n_1 - n_2$ particles sit on the edges, so in total we meet at most $n + 6$ edges.

We will use the sink comparison argument to compute $E^\infty_{2,0}$, see Section 2.3. By Corollary 2.6, we have

$$E^1_{p,q} \cong \bigoplus_{q_v+q_w=q} \bigoplus_{j\in J} H_{q_v}(\mathrm{Conf}_{S^j_v}(V_v)) \otimes H_{q_w}(\mathrm{Conf}_{S^j_w}(V_w))$$

for some finite index set J and $S^j_v \sqcup S^j_w \subset \underline{\mathbf{n}}$. The horizontal boundary map d_1 does not change the degrees q_v and q_w, so we can split each row of the E^1-page into chain complexes $\mathcal{C}^{q_v,q_w}$ indexed by tuples $(q_v, q_w) \in \mathbb{N}^2$. Again, we write $\mathcal{C}^{q_v,q_w}[Z]$ for the corresponding chain complexes of $E^1_{\bullet,\bullet}[Z]$.

Each such chain complexes with $q_v = 0$ or $q_w = 0$ stays the same when turning v or w into a sink, respectively. In particular, the bottom row is independent of Z: we have

$$E^2_{p,0}[Z] \cong E^2_{p,0}[\{v,w\}] \cong H_p(\mathrm{Conf}_n(B_k, \{v,w\}))$$

for any subset $Z \subset \{v,w\}$, where the last isomorphism follows since for $Z = \{v,w\}$ only the zeroth row of the first page is non-trivial.

We now use this general idea to describe $E^\infty_{2,0}$ and $E^\infty_{1,1}$. The module $E^\infty_{2,0} = E^3_{2,0}$ is given by the kernel of the map

$$d_2[\varnothing] \colon E^2_{2,0}[\varnothing] \to E^2_{0,1}[\varnothing].$$

By the comparison of E^2-pages as described above, we have

$$E^2_{\bullet,0}[\varnothing] \cong E^2_{\bullet,0}[\{v\}],$$

and under this identification the kernel of $d_2[\varnothing]$ is a submodule of the kernel of $d_2[\{v\}]$: the inclusion $\mathrm{Conf}_n(B_k) \hookrightarrow \mathrm{Conf}_n(B_k, \{v\})$ induces the canonical projection

$$E^2_{0,1}[\varnothing] \cong \mathcal{C}^{0,1} \oplus \mathcal{C}^{1,0} \twoheadrightarrow \mathcal{C}^{0,1} \cong E^2_{0,1}[\{v\}].$$

As $\mathrm{Conf}_n(B_k, \{v\})$ has only one essential vertex and no edge connecting two sinks, by Proposition 1.20 its combinatorial model is 1-dimensional and its second homology is zero, so this kernel and therefore $E^\infty_{2,0}$ are trivial.

It remains to construct generators for $E^\infty_{1,1} = E^3_{1,1}$. The map d_2 sends $E^2_{1,1}$ to zero, so it is sufficient to produce generators for

$$E^2_{1,1} \cong \mathcal{C}^{0,1}_1 \oplus \mathcal{C}^{1,0}_1 \cong \mathcal{C}^{0,1}_1[\{v\}] \oplus \mathcal{C}^{1,0}_1[\{w\}] \cong E^2_{1,1}[\{v\}] \oplus E^2_{1,1}[\{w\}].$$

By symmetry, we can further reduce the problem to finding generators for $E^2_{1,1}[\{v\}]$.

By the description above, $E^1_{p,q}[\{v\}]$ is non-trivial only for $q \leqslant 1$ and the space $\mathrm{Conf}_n(B_k, \{v\})$ is homotopy equivalent to a 1-dimensional complex. Therefore, $E^\infty_{1,1}[\{v\}] = E^3_{1,1}[\{v\}] = 0$ and

$$d_2[\{v\}] \colon E^2_{3,0}[\{v\}] \to E^2_{1,1}[\{v\}]$$

is surjective. The description of the bottom row above gives

$$E^2_{3,0}[\{v\}] \cong E^2_{3,0}[\{v, w\}] \cong H_3(\mathrm{Conf}_n(B_k, \{v, w\})).$$

By Proposition 4.14, the module $E^2_{3,0}[\{v\}]$ is generated by classes meeting at most $8 \leqslant n + 6$ edges. The corresponding classes in the zeroth homology groups can be represented by configurations without particles on the vertices v and w by moving all those particles onto one of the eight edges. Mapping such a class via d_2 to get the image in $E^2_{1,1}[\{v\}]$ (see Section 2.1.1) can be done without using more than those eight edges: the lifts we need to choose connect different 0-cycles involving only eight of the edges via one of the vertices. All such cycles can be connected by using only those edges because the configuration space of Star_8 is connected.

By the same argument, lifting the corresponding class to the zeroth column to get a representative in $\mathrm{Conf}_n(B_k)$ can be done in a way that only $8 \leqslant n + 6$ edges are met.

Therefore, the second homology of $\mathrm{Conf}_n(B_k)$ can be generated by classes involving at most $n + 6$ edges, which proves the claim. □

It remains to show the following two results:

Proposition 4.13. *For each* $k \in \mathbb{N}$ *the group* $H_2(\mathrm{Conf}_n(B_k, \{v, w\}))$ *is generated by classes meeting at most* 4 *edges.*

Proposition 4.14. *For each* $k \in \mathbb{N}$ *the group* $H_3(\mathrm{Conf}_n(B_k, \{v, w\}))$ *is generated by classes meeting at most* 8 *edges.*

Both proofs proceed by investigating the Mayer-Vietoris spectral sequence for the following open cover $\mathcal{V}$ of B_k: for each edge $e \in E(B_k)$, let V_e be a small contractible open neighborhood of the edge e. Denote by $E^*_{\bullet,\bullet}[Z]$ the Mayer-Vietoris spectral sequence for $\mathrm{Conf}_n(B_k, Z \cup \{v, w\})$ with $Z \subset B_k$, corresponding to the open cover $\mathcal{V}^{\underline{n}}$, see Section 2.1. Notice that we now allow the sink set Z to contain edges, and the two vertices are always sinks.

We will investigate the infinity page of the spectral sequence for $\mathrm{Conf}_n(B_k, \{v, w\})$

$$E^*_{\bullet,\bullet} = E^*_{\bullet,\bullet}[\varnothing] = E^*_{\bullet,\bullet}[\{v, w\}].$$

The main technique will be to turn whole edges into sinks and to compare the corresponding spectral sequences.

Remark 4.15. The fact that all vertices are sinks has the following implication: lifting elements of the infinity page to the zeroth column to interpret them as classes of $H_*(\mathrm{Conf}_n(B_k, Z))$ can be done without using any additional edges. In each step (see Section 2.1.1), the lifting along d_0 is given by choosing a k-chain in the star of one vertex whose boundary is the difference of two $(k-1)$-cycles. Since the configuration space of a star with any number of leaves such that only the central vertex is a sink is contractible, this chain can be chosen to not meet any new edges.

Therefore, it suffices to show that we can generate the modules on the infinity page by classes meeting only the claimed number of edges. △

4.3.1 The bottom row of the second page

Before going into the proofs, let us inspect the bottom row of the E^2-page of this spectral sequence. By the sink comparison argument again (see Section 2.3), the bottom row $E^2_{\bullet,0}$ is isomorphic to the bottom row $E^2_{\bullet,0}[Z]$ for $\mathrm{Conf}_n(B_k, Z)$ for *any* subset $\{v, w\} \subset Z \subset B_k$. In particular, it is the same for $Z = B_k$, for which we have $\mathrm{Conf}_n(B_k, B_k) = B_k^n$. The intersections of the open sets consist only of contractible connected components in this case, so all rows except the zeroth row are trivial already on the E^1-page. Therefore, the bottom row of the E^2-page is for any Z given by

$$E^2_{p,0}[Z] \cong E^2_{p,0}[B_k] \cong E^\infty_{p,0}[B_k] \cong H_p(B_k^n).$$

The homology group $H_p(B_k^n)$ is generated by products of p particles moving along embedded circles in B_k, without restrictions regarding collision. Using the relations in $H_1(B_k)$ we can choose generating systems of such products that have nice intersection patterns:

Remark 4.16. For $k \geqslant 3$ every circle given by edges e_1e_2 can be written as a difference of two circles e_1e_3 and e_2e_3, see Figure 4.1. If we now have a product of particle 1 and 2 both moving along the circle e_1e_2, then we can write this class of $H_2(B_k^n)$ as the difference of the two classes given by moving 1 along e_1e_2 and 2 along e_1e_3 or e_2e_3. Therefore, this product can be written as a sum of products where the circles share only one edge instead of two.

More generally, using the same idea we can find for $k \geqslant 2p-1$ a generating system of $H_p(B_k^n)$ consisting of such products such that

- each of the circles shares at most one of its two edges with any other circle, and
- all non-moving particles are on the sink v.

The restriction on p ensures that for each generator where the intersection of the circle factors is not as described above, there is an edge that is not contained in any of the embedded S^1, allowing us to resolve one of the intersections by writing a circle as the sum of two circles involving this free edge.

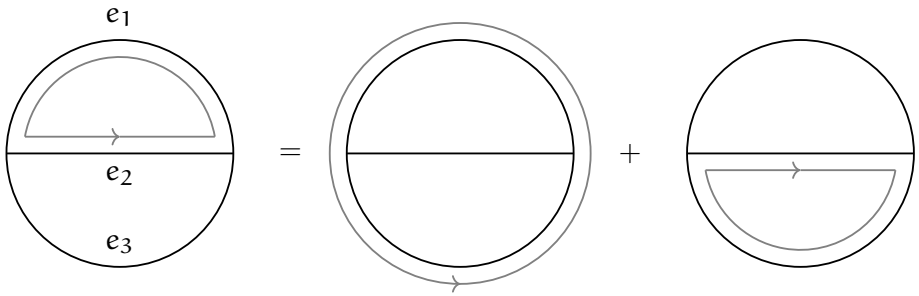

Figure 4.1: Replacing the movement of one particle along the circle e_1e_2 by two movements of this particle along e_1e_3 and e_2e_3.

In particular, at most $\lfloor \frac{p}{2} \rfloor$ edges are contained in two or more circles of one such product. △

We now describe how these product classes are represented in the chain complex $E^1_{\bullet,0}[B_k]$, see Section 2.1.1 for the description of how to interpret elements of the infinity page as classes in $H_*(\mathrm{Conf}_n(B_k, Z))$.

The movement of a particle $j \in \underline{\mathbf{n}}$ along an embedded circle given by two edges $e \neq e'$ is represented in $E^2_{1,0}[B_k]$ as the difference of the 0-classes given by putting j onto the two sinks in the intersection $V_e \cap V_{e'}$. Forgetting V_e from the intersection yields the trivial zeroth homology class of $V_{e'}$, and the lift connecting the two 0-cycles is given by moving j along e' from v to w (or from w to v). Forgetting $V_{e'}$ leads to the analogous lifting along the edge e, and combined they precisely give the 1-cycle moving j along the circle ee', see Figure 2.2.

The product of two S^1 movements along pairs of edges e_1e_1' and e_2e_2' is represented as follows: for easier description assume we only have two particles and particle j moves along the sequence of edges e_je_j' for $j = 1, 2$. To give the representation of this class in $E^2_{2,0}[B_k]$ we have to give classes in the zeroth homology of the intersection of three open sets. There are four open sets U_ϕ involving the edges e_j and e_j': one for each choice of $\phi(j) \in \{e_j, e_j'\}$ for $j = 1, 2$. In the nerve of the open cover of the configuration space there is a 2-dimensional cube on those four vertices, realized as the sum of two triangles. In both intersections corresponding to those triangles particle j is fixed in the intersection $V_{e_j} \cap V_{e_j'}$, so we can form the tensor product of the classes representing the individual S^1 movements as described above, see Figure 4.2. Mapping this sum via d_1 cancels the diagonal in the interior of the cube, and in each face on the boundary of the cube one of the two particles is freed. Hence, this maps to zero and lifting the class to the zeroth column shows that it realizes the product of the two S^1 movements as desired, see Section 2.1.1. If we have more than those two particles, we put all remaining particles onto the sink v in the open neighborhood of the first edge of B_k (for some arbitrary ordering of the edges).

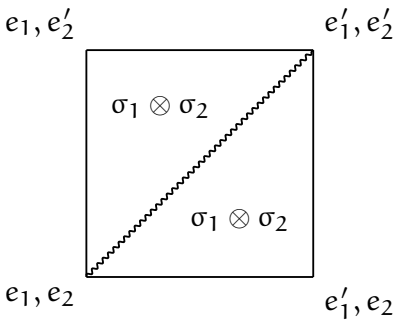

Figure 4.2: Representing a product of S^1-movements in $E^2_{2,0}$. The rectangle lives in the nerve of the open cover, the labels at each vertex indicate that the corresponding vertex is the open set U_ϕ for ϕ determined by the label $\phi(1), \phi(2)$. In both triangles, the particles 1 and 2 are restricted to the intersection of the open neighborhoods of their respective edges, and we take the class $\sigma_1 \otimes \sigma_2$ given by the tensor product of the classes representing the S^1 movements of the individual particles. The wiggly line cancels under d_1, the class $\sigma_1 \otimes \sigma_2$ restricted to the remaining faces represents zero because one of the two particles is freed.

Higher products of k movements along embedded S^1 are similarly presented by choosing a triangulation of a cube on 2^k vertices in the nerve of the open cover.

Given a product as described above in $E^2_{p,0}[Z]$ for $p \geqslant 2$, we can look at its image under $d_2[Z]$. We will be most interested in the situation where Z is the union of some of the closed edges of B_k, so we will only describe those images in such a situation. The map $d_2[Z]$ is constructed by choosing a representative in $E^0_{p,0}[Z]$, mapping it once via $d_1[Z]$, choosing a preimage under $d_0[Z]$ and again mapping it under $d_1[Z]$. The first map $d_1[Z]$ either frees precisely one of the particles or keeps all of them fixed. The latter parts have to sum to zero (already on the zeroth page), so they lift to zero under $d_0[Z]$, meaning that we only care about the parts where one of the particles is freed. The lifting along $d_0[Z]$ then is given by choosing a path of the freed particle between v and w. Mapping the result again via $d_1[Z]$ possibly frees another particle in each of the summands. If the two freed particles are freed to different edges or if the edge they are freed to is contained in Z, then this maps to zero. If they are freed to the same non-sink edge, then there are four non-trivial summands adding up to an H-class of the two particles on e. Recall that homology classes of $\mathrm{Conf}_S([0,1],\{0,1\})$ are also called H-classes, see Remark 3.4.

The higher boundary maps are defined similarly. The product of q H-classes on different edges and p circle classes is represented in $E^1_{p,q}$ in the same way: the H-classes are given in the tensor factors $H_1(\mathrm{Conf}_{S^j_e}(e,\{v,w\}))$, and the circle classes

are represented as above.

4.3.2 H-classes are stably trivial

The general idea for the description of generators for $H_*(\mathrm{Conf}_n(B_k, \{v, w\}))$ is that H-classes are stably trivial, meaning that for $k \gg q$, the product of q H-classes is trivial. For banana graphs with many edges we use this idea to show that the homology is assembled from circle classes. The following propositions make these ideas precise and will be the main ingredients for proving Proposition 4.13 and Proposition 4.14. For brevity, for a set of edges $\xi \subset E(B_k)$ we write

$$B_k - \xi = B_k - \bigcup_{e \in \xi} \mathring{e}.$$

Proposition 4.17. *For $\xi \subsetneq E(B_k)$ and each $q > 0$ we have $E^{\infty}_{p,q}[B_k - \xi] = 0$ for all p.*

Proof. Since $\xi \neq E(B_k)$ there exists at least one sink edge, and slowly collapsing this edge to a point shaped sink the space $\mathrm{Conf}_n(B_k, B_k - \xi)$ is homotopy equivalent to the configuration space of a (pointed) wedge of circles where all but $|\xi|$ of the circles are sinks, see Figure 4.3 and Corollary 1.8. By Proposition 3.12, the homology of the latter space is generated by products of configurations in the individual wedge summands.

Hence, the homology of $\mathrm{Conf}_n(B_k, B_k - \xi)$ is generated by products of particles moving along embedded circles that do not intersect in the interiors of the edges in ξ. These products are all realized in the bottom row $E^{\infty}_{\bullet,0}[B_k - \xi]$ by the description in Section 4.3.1, showing that $E^{\infty}_{p,q}[B_k - \xi] = 0$ for all $q > 0$. □

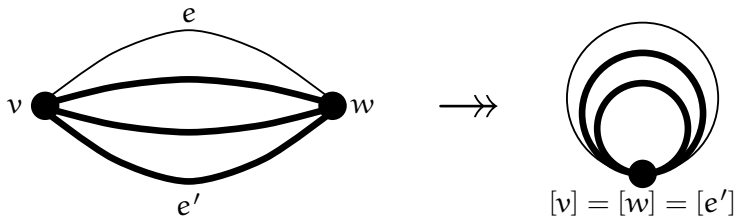

Figure 4.3: Collapsing a sink edge e' induces a homotopy equivalence on configuration spaces. The homology of the configuration space of the collapsed space is generated by products of particles moving along individual S^1 wedge summands such that no non-sink circle is used by more than one particle.

Proposition 4.18. *Let $q > 0$. Every product of* H*-classes and circle classes in $E^2_{p,q}[B_k - \xi]$ such that*

- *each edge met by at least two factors of the product is a sink edge, and*
- *we either have* $\xi \neq E(B_k)$ *or there are at least two edges not met by the product*

is contained in the image of d_2*. In particular, for all* $k \geqslant q+2$ *we have* $E^{\infty}_{0,q} = E^{3}_{0,q} = 0$*.*

Proof. The product of two circle classes sharing precisely one non-sink edge e maps under d_2 to the H-class of the two particles on e. Assume that we have a product in $E^2_{p,q}[B_k - \xi]$ as described in the statement with an H-class of particles x and x' on the edge e. We can replace this H-class by the product of x and x' moving along circles intersecting in e and possibly one sink edge, see Figure 4.4. For this, we used that there is either a sink edge or two non-sink edges not met by the product. This determines an element of $E^2_{p+2,q-1}[B_k - \xi]$. By the conditions on the product, this maps under d_2 to the product we started with.

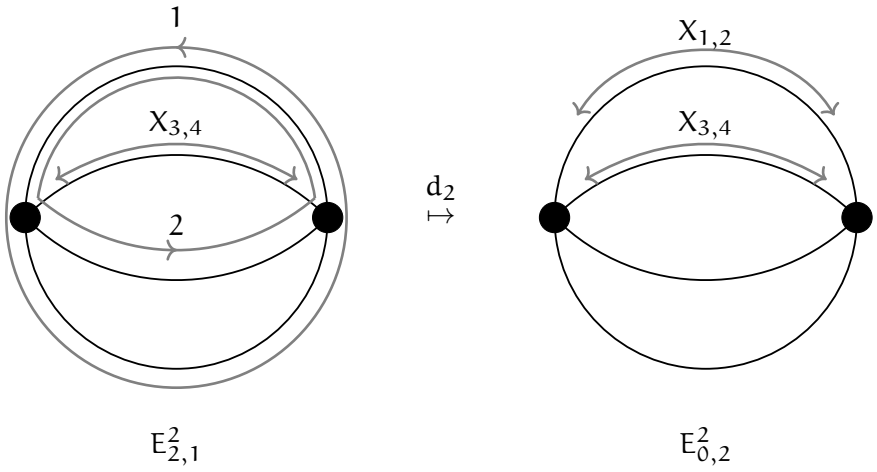

Figure 4.4: Constructing a d_2-preimage of a product of two H-classes $X_{1,2}$ and $X_{3,4}$ involving particles $\{1,2\}$ and $\{3,4\}$, respectively. We use two free edges in order to form a product of circle classes of the particles 1 and 2 intersecting in precisely one edge.

Since every H-class can be written as a sum of H-classes involving only two particles by Proposition 1.24, this proves the claim. □

Proposition 4.19. *Let* $\Sigma_{i\in I}\alpha_i X_i \in E^2_{p,q}[B_k - \xi]$ *such that each* X_i *is a product of* H*-classes and circle classes.*
(i) Let for $\{e_1,\dots,e_q\} \subset E(B_k)$ *the index set* $I_{\{e_1,\dots,e_q\}} \subset I$ *be given by those indices* i *such that* X_i *has* H*-factors on each of those edges. Then* $\Sigma_{i\in I}\alpha_i X_i = 0$ *implies* $\Sigma_{i\in I_{\{e_1,\dots,e_q\}}}\alpha_i X_i = 0$*.*

(ii) If each H*-class of each* X_i *involves precisely two particles, then we have the following: for* $S \subset \underline{\mathbf{n}}$ *of cardinality* $2q+p$ *denote by* I_S *the index set of all* i *such that the set of non-moving particles of* X_i *is given by* $\underline{\mathbf{n}} - S$*, then* $\Sigma_{i\in I}\alpha_i X_i = 0$ *implies* $\Sigma_{i\in I_S}\alpha_i X_i = 0$*.*

Proof. The first claim follows immediately from the decomposition

$$E^1_{p,q}[B_k - \xi] \cong \bigoplus_{\Sigma q_e = q} \bigoplus_{j\in J} \bigotimes_e H_{q_e}(\mathrm{Conf}_{S^j_e}(e, Z_e))$$

as described in Corollary 2.6, and the fact that d_1 preserves the homology degrees q_e. Here, $Z_e = \{v, w\}$ for $e \in \xi$ and $Z_e = e$ otherwise.

We now prove the second claim. Recall that all fixed particles are fixed on the sink v. The projection map π_S and the one-sided inverse given by adding the particles $\underline{\mathbf{n}} - S$ to the sink v map the open sets onto each other in an order preserving way because we ordered the open sets lexicographically. By a slight variation of Remark 2.4 they induce maps of Mayer-Vietoris spectral sequences, sending elements which would land in a degenerate simplex of the nerve of the open cover to zero. The variation is the following: for the inverse map we have a choice in which open set we put the fixed particles on v instead of having one canonical choice, and we just fix an arbitrary edge and put all fixed particles into the open set containing that edge.

Given a linear combination as in the statement, the map induced by the projection π_S to the particles S kills all X_i where the set of moving particles is not precisely given by S: if one of the particles of $\underline{\mathbf{n}} - S$ is part of a circle class then the simplices of the 2^p-cube in the nerve of the open cover map to degenerate simplices, so they map to zero. If one of those particles is part of an H-class then the corresponding H-class maps to zero because $\mathrm{Conf}_1([0,1],\{0,1\}) \simeq [0,1]$ is contractible.

The image of $\Sigma_{e\in I}\alpha_i X_i$ under the induced map on the E^2-page of the composition of those two maps is therefore precisely given by $\Sigma_{e\in I_S}\alpha_i X_i$. Since this composition of maps also induces a map on the E^3-page, this is an element of the kernel of d_2 as claimed. □

4.3.3 Generators for banana graphs with sinks

We are now ready to give generators for $H_2(\mathrm{Conf}_n(B_k, \{v, w\}))$.

Proof of Proposition 4.13. Assume $k \geqslant 5$ since for $k \leqslant 4$ there is nothing to prove.

We will show that the modules $E^\infty_{p,q}$ for $p+q=2$ are generated by classes involving at most 4 edges of B_k. By Remark 4.15, this is enough to produce a generating system for $\mathrm{Conf}_n(B_k, \{v, w\})$ meeting at most 4 edges.

The module $E^\infty_{0,2}$ is trivial by Proposition 4.18.

The module $E^{\infty}_{1,1}$ is a quotient of $E^2_{1,1}$, so we give generators for

$$E^2_{1,1} \cong \bigoplus_{e \in E(B_k)} E^2_{1,1}[B_k - e].$$

By Proposition 4.17, we have for each $e_0 \in E(B_k)$

$$E^3_{1,1}[B_k - e_0] = E^{\infty}_{1,1}[B_k - e_0] = 0,$$

and therefore that

$$d_2[B_k - e_0] \colon E^2_{3,0}[B_k - e_0] \to E^2_{1,1}[B_k - e_0]$$

is surjective. The module

$$E^2_{3,0}[B_k - e_0] \cong E^2_{3,0}[B_k] \cong H_3(B^n_k)$$

is generated by products of three particles moving along embedded circles. As described in Remark 4.16, we can choose a generating system consisting of such products where the circles only intersect in one edge. If the circles meet in an edge that is not e_0, such a product maps to zero under $d_2[B_k - e_0]$. The remaining products map to something non-trivial, but because they do not intersect in any other edge, the image of the corresponding element under

$$d_2 \colon E^2_{3,0} \to E^2_{1,1} \cong \bigoplus_{e \in E(B_k)} E^2_{1,1}[B_k - e]$$

is zero in all direct summands $E^2_{1,1}[B_k - e]$ for $e \neq e_0$, and the following square commutes for each $e_0 \in E(B_k)$:

$$\begin{array}{ccc} E^2_{3,0} & \xrightarrow{\quad d_2 \quad} & E^2_{1,1} \\ \uparrow & & \uparrow \\ E^2_{3,0}[B_k - e_0] & \xrightarrow{d_2} & E^2_{1,1}[B_k - e_0] \end{array}$$

This implies that $d_2 = d_2[\varnothing]$ is surjective onto each of its direct summands $E^2_{1,1}[B_k - e_0]$, so we have $E^{\infty}_{1,1} = 0$.

The module $E^{\infty}_{2,0}$ is the kernel of the map

$$d_2 \colon E^2_{2,0} \to E^2_{0,1}.$$

The source of this map is given by $H_2(B^n_k)$, so it is generated by products of two particles following embedded circles. By Remark 4.16 again, we can generate $H_2(B^n_k)$

by products where the embedded circles share at most one edge. Products along circles only meeting at the vertices are elements of $E^{\infty}_{2,0}$ meeting four edges, so it remains to investigate linear combinations of products sharing precisely one edge.

Let such a linear combination $\Sigma\alpha_i X_i$ in the kernel of d_2 be given. Each X_i maps under d_2 to an H-class, so by Proposition 4.19, we can assume that $n = 2$ and that the two circles of each X_i intersect in the same edge e_0. Changing one of the edges which is only met by one of the two circles does not change the image under d_2, so any choice of doing that produces a kernel element meeting four edges. Adding these kernel elements we can arrange that each of the two particles moves for all X_i along the same pair of edges (one of which is e_0). This means that the X_i are (up to a sign) all the same class, and since that class does not map to zero under d_2 we have reached the trivial linear combination.

This shows that the second homology of $\mathrm{Conf}_n(B_k, \{v, w\})$ is generated by classes meeting at most four edges. □

We will now prove Proposition 4.14. This proof works along the same lines as the previous one, so we omit a few details. Recall the definition of the Mayer-Vietoris spectral sequences $E^*_{\bullet,\bullet}[Z]$ for $\mathrm{Conf}_n(B_k, Z \cup \{v, w\})$ with $Z \subset B_k$ a union of edges. The open cover of B_k used in the construction is given by contractible neighborhoods of the edges.

Proposition 4.20. *The module $E^{\infty}_{0,3}$ is trivial for $k \geqslant 5$.*

Proposition 4.21. *The module $E^{\infty}_{2,1}$ is trivial for $k \geqslant 7$.*

Proposition 4.22. *The module $E^{\infty}_{1,2}$ is trivial for $k \geqslant 9$.*

Proposition 4.23. *The module $E^{\infty}_{3,0}$ is generated by classes meeting at most 6 edges.*

Remark 4.24. Despite some effort, we were not able to show that $E^{\infty}_{p,q}$ is trivial for all $k \gg q > 0$. We were able to prove this for a few more cases, but since we are not aware of any implications of such a result for ordinary configuration spaces, we decided not to include those rather technical arguments. △

Proof of Proposition 4.14. By Remark 4.15, this follows from the four propositions above. □

We will now prove the four propositions above.

Proof of Proposition 4.20 ($E^{\infty}_{0,3}$). This follows from Proposition 4.18. □

Proof of Proposition 4.21 ($E^{\infty}_{2,1}$). Assume that we have $k \geqslant 7$, because otherwise there is nothing to show. We need to show that $E^{\infty}_{2,1} = E^{3}_{2,1} = 0$. As in the proof of Theorem G, we have

$$E^{2}_{2,1} = E^{2}_{2,1}[\{v, w\}] \cong \bigoplus_{e \in E(B_k)} E^{2}_{2,1}[B_k - e].$$

By Proposition 4.17, the map $d_2[B_k - e]$ landing at position $(2,1)$ has to be surjective, and a generating system for $E^{2}_{4,0}[B_k - e]$ determines a generating system for $E^{2}_{2,1}[B_k - e]$ via $d_2[B_k - e]$.

By Remark 4.16, we can choose for fixed e_0 a generating system for $E^{2}_{4,0}[B_k - e_0] \cong H_4(B^n_k)$ given by products of four particles moving along embedded circles in B_k meeting in at most two edges.

Amongst these generators there are some where at most one of the four embedded circles meets e_0. These classes map to zero under $d_2[B_k - e_0]$, so to produce a generating system for $E^{2}_{2,1}[B_k - e_0]$ we can ignore them. The remaining generators now come in two different types: For the first type of generator, the edge e_0 is contained in at least two circles and all other edges are contained in at most one circle. For the second type, e_0 and precisely one additional edge e_1 are each contained in precisely two circles. For a schematic picture of the two types see Figure 4.5.

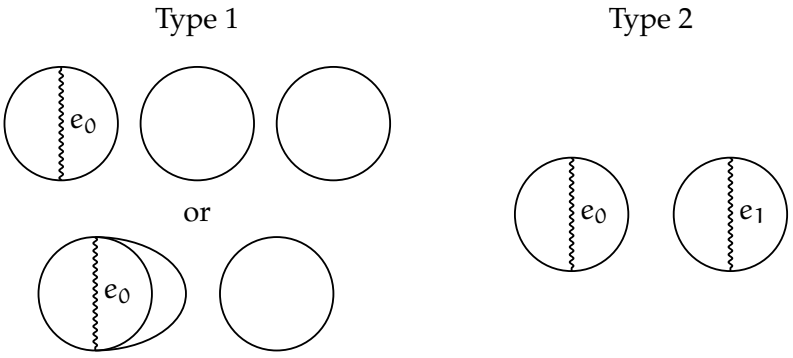

Figure 4.5: A schematic picture of the two types of products of particles moving along circles. A wiggly edge is contained in multiple circles, the other edges are contained in exactly one circle. Note that this is only schematic, in reality the circles all meet at the vertices v and w.

These two types do not map to zero under $d_2[B_k - e_0]$, so their images represent generators for $E^{2}_{2,1}[B_k - e_0]$. This describes a generating set for $E^{2}_{2,1} \cong \bigoplus_e E^{2}_{2,1}[B_k - e_0]$, and we will now use that description to show $E^{3}_{2,1} = 0$.

Each generator given by the image of a product of the **first type** is also in the image

of

$$d_2 \colon E^2_{4,0}[\{v,w\}] \to E^2_{2,1}[\{v,w\}]$$

because the corresponding products sharing only one edge are mapped in precisely the same fashion as in $E^2_{\bullet,\bullet}[B_k - e_0]$ for the corresponding e_0. Therefore, they represent zero in $E^3_{2,1}[\{v,w\}]$ and we only have to consider generators coming from products of the second type.

A generator coming from a product of the **second type** is a product of an H-class and two circles sharing another edge. Each such generator maps under d_2 to a product of two H-classes at these edges. Hence, given a linear combination $\Sigma\alpha_i X_i$ of such generators in the kernel of d_2, we can by Proposition 4.19 assume that these two edges are the same for all X_i, and we call them e_0 and e_1.

The image of a product of four circles X of the second type maps under

$$d_2 \colon E^2_{4,0}[\{v,w\}] \to E^2_{2,1}[\{v,w\}] \cong \bigoplus_{e \in E(B_k)} E^2_{2,1}[B_k - e]$$

to $d_2[B_k - e_0](X) + d_2[B_k - e_1](X)$, so by adding $d_2[\{v,w\}]$-boundaries we can assume that all X_i are in $E^2_{2,1}[B_k - e_0]$.

Given one fixed X_i where the two circle classes have free edges e_2 and e_3 we can look at X_i', defined as X_i with e_2 replaced by an edge e_2' which is distinct from the other four edges. The difference $X_i - X_i'$ is (up to d_1-boundaries) the product of an H-class at e_0 and two disjoint circle classes along $e_2 e_2'$ and $e_1 e_3$. By the fact that we have at least seven edges and Proposition 4.18, it is in the image of d_2 and represents the trivial class on the third page.

Adding those d_2-boundaries we can arrange that all X_i are the same, and since $d_2(X_i) \neq 0$, we have reached the trivial linear combination. This shows that $E^\infty_{2,1} = E^3_{2,1} = 0$. □

Proof of Proposition 4.22 ($E^\infty_{1,2}$). Assume $k \geqslant 9$ because otherwise, there is nothing to show. We will show that $E^\ell_{1,2}$ is a quotient of

$$\bigoplus_{e_1 \neq e_2} E^\ell_{1,2}\big[B_k^{e_1 e_2}\big]$$

for all $\ell \geqslant 2$, where we write $B_k^{e_1 e_2} := B_k - \mathring{e}_1 - \mathring{e}_2$. Since by Proposition 4.17 we have $E^4_{1,2}\big[B_k^{e_1 e_2}\big] = 0$ for all choices of $e_1 \neq e_2$, this will prove that $E^\infty_{1,2} = 0$.

For $\ell = 2$ this is true (even without the quotient):

$$E^2_{1,2} = E^2_{1,2}[\{v,w\}] \cong \bigoplus_{e_1 \neq e_2 \in E(B_k)} E^2_{1,2}\big[B_k^{e_1 e_2}\big],$$

so we need to show the claim for $\ell = 3$ and $\ell = 4$.

As a first step, we will investigate the image of

$$d_2 \colon E^2_{3,1}\left[B_k^{e_1 e_2}\right] \cong E^2_{3,1}\left[B_k^{e_1}\right] \oplus E^2_{3,1}\left[B_k^{e_2}\right] \to E^2_{1,2}\left[B_k^{e_1 e_2}\right].$$

By Proposition 4.17 again, we see that

$$d_2 \colon E^2_{5,0}[B_k^e] \to E^2_{3,1}[B_k^e]$$

is surjective for each e, and we can choose a generating system of $E^2_{3,1}\left[B_k^e\right]$ by mapping a generating system of $E^2_{5,0}\left[B_k^e\right]$ via this map. We choose this system such that at most two edges are contained in two or more circles, see Remark 4.16. Here, we used that we have at least nine edges. The corresponding elements of $E^2_{3,1}\left[B_k^{e_1}\right]$ where at most *one* circle moves along e_2 are also in the image of

$$d_2 \colon E^2_{5,0}\left[B_k^{e_1 e_2}\right] \to E^2_{3,1}\left[B_k^{e_1 e_2}\right],$$

so they map trivially to $E^2_{1,2}\left[B_k^{e_1 e_2}\right]$. The remaining classes have at least two circles moving along e_2, so by the choice of basis the circles do not intersect in any edge other than e_1 and e_2. Therefore, they map in precisely the same way under

$$d_2 \colon E^2_{3,1}[\{v,w\}] \to E^2_{1,2}[\{v,w\}].$$

Repeating this argument for all choices of edges e_1 and e_2 we get that $E^3_{1,2}[\{v,w\}]$ is a quotient of

$$\bigoplus_{e_1 \neq e_2} E^3_{1,2}\left[B_k^{e_1 e_2}\right].$$

Therefore, it suffices to realize all elements in the image of the surjective map

$$d_3 \colon E^3_{4,0}\left[B_k^{e_1 e_2}\right] \to E^3_{1,2}\left[B_k^{e_1 e_2}\right]$$

as images under the map

$$d_3 \colon E^3_{4,0}[\{v,w\}] \to E^3_{1,2}[\{v,w\}].$$

Each element in $E^3_{4,0}\left[B_k^{e_1 e_2}\right] = \ker\left(E^2_{4,0}\left[B_k^{e_1 e_2}\right] \to E^2_{2,1}\left[B_k^{e_1 e_2}\right]\right)$ can be written as a sum of generators of $E^2_{4,0}\left[B_k^{e_1 e_2}\right]$. These generators can be taken to be products of circle classes meeting in at most two edges by Remark 4.16, which is possible since we have at least 9 edges. There are four types of such products, differing in the intersection pattern of the four embedded circles:

- two circles share one of the two edges e_1 and e_2 and the other edge is contained in at most one circle,
- two circles share e_1 and two circles share e_2,

- three circles share one of the two edges and
- all four circles share one of the two edges.

For a schematic representation of these four types see Figure 4.6. It remains to show that each element of the kernel can be written as a linear combination of these products such that there is no product of the second kind of type 1: this is enough because the remaining classes do not intersect outside of the edges e_1 and e_2, so these classes of $E^3_{4,0}\left[B_k^{e_1e_2}\right] \cong E^3_{4,0}[\{v,w\}]$ map in the same way under

$$d_3 \colon E^3_{4,0}[\{v,w\}] \to E^3_{1,2}[\{v,w\}],$$

showing that $E^4_{1,2}[\{v,w\}]$ is a quotient of the trivial module

$$\bigoplus_{e_1 \neq e_2} E^4_{1,2}\left[B_k^{e_1e_2}\right] = \bigoplus_{e_1 \neq e_2} E^\infty_{1,2}\left[B_k^{e_1e_2}\right] = 0.$$

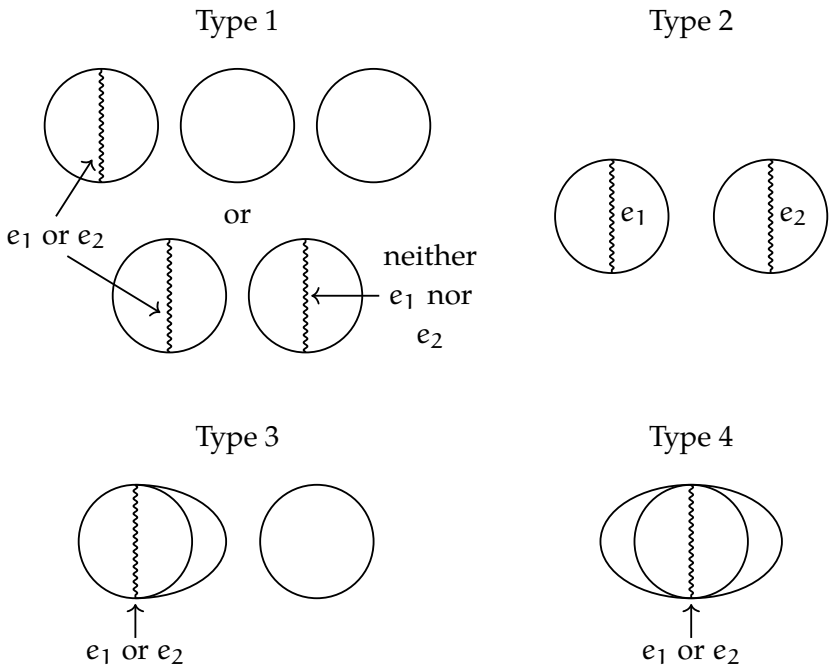

Figure 4.6: A schematic picture of the four types of products of movements of particles along circles. A wiggly edge is contained in multiple circles, the other edges are contained in exactly one circle. Note that this is only schematic, in reality the circles all meet at the vertices v and w.

Now assume we have a linear combination $X = \Sigma \alpha_i X_i$ of products of these types that maps to zero under d_2. Our claim is that none of the X_i is given by a product of the second kind of type 1. By Proposition 4.19, we can assume $n = 4$.

All X_i of the second kind of type 1 with circles sharing e_1 and one other fixed edge e can be modified in the following way: by adding products of the first kind of type 1 we can change the free edges of the two circles sharing the edge e, see Remark 4.16. Since we have more than seven edges, we can thus arrange that for each edge $e \notin \{e_1, e_2\}$ and each pair of particles $x \neq x'$ there is at most one X_i of the second kind of type 1 consisting of circles sharing e_1 and e such that x and x' meet e. By the same argument, we can assume the same for e_2 instead of e_1.

Assume that X_{i_0} is such a product involving (after renaming) the particles $\{1,2\}$ moving along e_1 and $\{3,4\}$ moving along $e \notin \{e_1, e_2\}$. We will now show that α_{i_0} has to be zero. Look at the map

$$\mathrm{Conf}_4(B_k, B_k^{e_1 e_2}) \hookrightarrow \mathrm{Conf}_{\{1,2\}}(B_k, B_k^{e_1}) \times \mathrm{Conf}_{\{3,4\}}(B_k, B_k)$$

given by the product of the projection maps $\pi_{\{1,2\}}$ and $\pi_{\{3,4\}}$. We can define the analogous open cover of the right hand side by open sets indexed by $\phi\colon \underline{4} \to E(B_k)$ restricting particle i to the open neighborhood of the edge $\phi(i)$. This defines a Mayer-Vietoris spectral sequence, and since the map above maps an open set $U_{\phi_0 \cdots \phi_p}$ by the inclusion into the open set in the product with that same label $\phi_0 \cdots \phi_p$, this inclusion induces a map of spectral sequences, see Remark 2.4. Denote by $\overline{X}_i$ the image of X_i in the spectral sequence of the product denoted by $\overline{E}^*_{\bullet,\bullet}$.

Each $\overline{X}_i$ where the circles of particles 1 and 2 do not intersect in e_1 are elements of the infinity page, so after removing all those elements from the linear combination the rest will still be in the kernel of d_2. Therefore, we assume that all $\overline{X}_i$ move 1 and 2 along e_1. We have for each $p \geqslant 0$

$$\begin{aligned}
\overline{E}^1_{p,1} &\cong \bigoplus_{\Phi_p} H_1(\overline{U}_{\Phi_p}) \\
&\cong \bigoplus_{\Phi_p} H_1\left(\overline{U}^{12}_{\Phi_p} \times \overline{U}^{3,4}_{\Phi_p}\right) \\
&\cong \bigoplus_{\Phi_p} H_1\left(\overline{U}^{12}_{\Phi_p}\right) \otimes H_0\left(\overline{U}^{34}_{\Phi_p}\right) \oplus H_0\left(\overline{U}^{12}_{\Phi_p}\right) \otimes H_1\left(\overline{U}^{34}_{\Phi_p}\right) \\
&\cong \bigoplus_{\Phi_p} H_1\left(\overline{U}^{12}_{\Phi_p}\right) \otimes H_0\left(\overline{U}^{34}_{\Phi_p}\right),
\end{aligned}$$

where the sum is over all triples $\Phi_p = \{\phi_0 < \cdots < \phi_p\}$, the open set $\overline{U}_{\Phi_p}$ is the open set in the product corresponding to Φ_p, and $\overline{U}^{xy}_{\Phi_p} := \pi_{\{x,y\}}(\overline{U}_{\Phi_p})$. The last line is true because all intersections of open sets in $\mathrm{Conf}_{\{3,4\}}(B_k, B_k)$ have contractible path components and therefore do not have first homology. The module $H_1(\overline{U}^{12}_{\phi_0 \cdots \phi_p})$ is only non-trivial if $\phi_i(\{1,2\}) = \{e_1\}$ for all $0 \leqslant i \leqslant p$, and then it is given by $H_1(\mathrm{Conf}_{\{1,2\}}(e, \{v, w\})) \cong \mathbb{Z}$. This shows that $\overline{E}^2_{\bullet,1}$ is isomorphic (as chain

complexes) to the bottom row of the analogous Mayer-Vietoris spectral sequence for $\mathrm{Conf}_{\{3,4\}}(B_k, B_k)$. Therefore, we have

$$\overline{E}^2_{2,1} \cong H_2\left(B_k^{\{3,4\}}\right) \cong H_1(B_k) \otimes H_1(B_k).$$

The $\overline{X}_i$ map under this identification to products of the particles 3 and 4 moving along embedded circles. By assumption, only $\overline{X}_{i_0}$ maps to such a product where both particles move along the fixed edge e, all other products are either disjoint or share an edge which is not e. In order to add to zero in $H_1(B_k) \otimes H_1(B_k)$, we therefore must have $\alpha_{i_0} = 0$.

Repeating this argument with all X_i of the second kind of type 1 shows that all these coefficients are zero, which proves the claim. □

Proof of Proposition 4.23 ($E^\infty_{3,0}$). Let $k \geqslant 7$, because otherwise there is nothing to show. By Proposition 4.18, we have $E^3_{0,2} = 0$, which implies $E^3_{3,0} = E^\infty_{3,0}$. It remains to give generators for the kernel of $d_2 \colon E^2_{3,0} \to E^2_{1,1}$. Since $k \geqslant 7 \geqslant 5$, we can choose a generating system of $E^2_{3,0} \cong H_3(B_k^n)$ such that the circles meet in at most one edge, see Remark 4.16. As described above, d_2 maps the products without shared edges to zero, the other X_i create H-classes at that shared edge. By Proposition 4.19, we can further assume $n = 3$ and that all intersecting circles intersect only in one fixed edge e_0.

Let $\Sigma_{i\in I}\alpha_i X_i \in E^2_{3,0}$ be a linear combination of such generators that maps to zero under d_2. The X_i mapping to zero meet only 6 edges, so we can assume there are none. For a schematic description of all such products sharing one edge, see Figure 4.7. Recall that an edge of one of the circles is called *free* if it is not shared with any other circle class in the product.

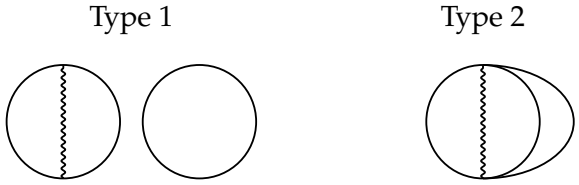

Figure 4.7: A schematic picture of products of type 1 and 2. The wiggly edge e_0 is met by multiple circles, the other edges are met by exactly one circle. Note that this is only schematic, in reality the circles all meet at the vertices v and w.

Choose three distinct edges e_1, e_2, e_3 such that none of them is equal to e_0. We can change the free edges of the three circles of each product of type 2 by adding type 1 products such that particle i moves along edge e_i for $1 \leqslant i \leqslant 3$, see Remark 4.16.

Here, we used that we have at least five edges. Therefore, all products of type 2 are the same (up to a sign), and by combining the corresponding α_i we can assume that there is precisely one X_{i_0} that is of type 2. We will now show that we must have $\alpha_{i_0} = 0$.

The image $d_2(\Sigma\alpha_i X_i)$ contains precisely one summand Z with an H-movement of the particles $\{1,2\}$ on e_0 and particle 3 moving along the circle e_0e_3. The coefficient of this product is $\pm\alpha_{i_0}$, and it is represented by an element in the direct summand $H_1(U_{\phi_1\phi_2}) \subset E^1_{1,1}$ for $\phi_1(\{1,2,3\}) = \phi_2(\{1,2\}) = \{e_0\}$ and $\phi_2(3) = e_3$. Now let $Y = \Sigma_{j\in J}\beta_j Y_j \in E^1_{2,1}$ be such that

$$d_2(\Sigma\alpha_i X_i) + d_1(Y) = 0 \in E^1_{1,1}.$$

The summand Z is the only one which is contained in a direct summand with particles 1 and 2 forming an H-class on e_0 and particle 3 moving along a circle involving e_0. Therefore, the element $d_1(Y)$ must be non-trivial in that direct summand. This summand is only hit by $d_1(Y_j)$ if Y_j is the given by the H-class of the particles 1 and 2 on e_0 with the third particle in the intersection of the neighborhoods of e_0, e_3 and one additional edge e^j. It is straightforward to see that by adding d_1-boundaries of the analogous classes where particle 3 is in the intersection of *four* instead of three edge neighborhoods, we can assume that

- for each Y_j given by an H-class of 1 and 2 on e_0 and particle 3 in the intersection of the neighborhoods of a triple of edges, two of these three edges are e_0 and e_3, and
- for each edge $e \in E(B_k)$ there is at most one such Y_j as above such that $e^j = e$.

The image $d_1(Y_j)$ of such a Y_j has a summand given by the product of an H-class of 1 and 2 on e_0 and particle 3 moving along e_0e^j. By the conditions above, Y_j is the only element landing in the direct summand of $E^1_{1,1}$ containing that element, and since $d_2(\Sigma\alpha_i X_i)$ does not hit that direct summand either we must have $\beta_j = 0$. Therefore, the direct summand of Z is not hit by d_1 and we must have $\alpha_{i_0} = 0$, so we can assume that all X_i in our linear combination are products of type 1.

To cancel in $E^2_{1,1}$, the products X_i of type 1 have to at least land in the same direct summand, so we can assume that all X_i map to elements where the H-class on e_0 involves the particles $\{1,2\}$. Now fix an edge e_3. Changing one of the free edges of one X_i gives a class X_i' with the same image image under d_2, so by adding such differences $X_i - X_i'$ we can again arrange that particles 1 and 2 never move along a circle involving e_3. The elements we added are kernel elements meeting at most six edges, so we can forget about them. Using Remark 4.16 again we can assume that for each X_i particle 3 moves along a circle e_3e^i for some edge e^i. Note also that none of the e^i is equal to e_0 since the X_i are of type 1.

Each $d_2(X_i)$ is now the product of an H-class on e_0 with particle 3 in the intersection of the neighborhoods of e_3 and e^i. Changing the free edges of particles 1 and 2 again and combining coefficients we can assume that each such product only comes from *one* X_i. It is straightforward to check that these elements of $E^2_{1,1}$ are linearly independent, so

$$d_2(\Sigma \alpha_i X_i) = \Sigma \alpha_i d_2(X_i) = 0$$

implies $\alpha_i = 0$ for all i.

This shows that $E^\infty_{3,0}$ is generated by classes meeting at most six edges. □

4.4 Stability for the second homology group

In this section we will prove the following theorem, showing stabilization for the second homology group. Recall the definition of $\mathbf{H}^{\underline{\Gamma}}_{q,n}$ from Section 4.2.

Theorem H. *For any choice of graphs* K_i, G_i *and* $n \in \mathbb{N}$*, the* $\mathbf{FI}^{\times \ell}$*-module* $\mathbf{H}^{\underline{\Gamma}}_{2,n}$ *is finitely generated in degree* $(n+6, \dots, n+6)$*.*

Given $\underline{\Gamma} = \{(K_0, G_0), \dots, (K_\ell, G_\ell)\}$ we now define a generalization of $\mathbf{H}^A_{q,n}\underline{\Gamma}$ including sinks. Let $Z_i \subset V(G_i)$ for $0 \leqslant i \leqslant \ell$ and write $\mathcal{Z} = (Z_0, \dots, Z_\ell)$. Given this data, define the following generalization of $\mathbf{H}^A_{q,n}\underline{\Gamma}$:

$$\mathbf{H}^A_{q,n}(\underline{\Gamma}, \mathcal{Z}) \coloneqq H_q(\mathrm{Conf}_n(\mathbf{G}_{\underline{\Gamma}}, Z_{\mathbf{G}_{\underline{\Gamma}}}); A),$$

where $Z_{\mathbf{G}_{\underline{\Gamma}}}$ is given by the union of all vertex sets Z_i in all copies of the G_i in $\mathbf{G}_{\underline{\Gamma}}$.

We will prove the following generalization of Theorem H.

Theorem 4.25. *For any choice of* $\underline{\Gamma}$ *and* $\mathcal{Z}$ *and each* $n \in \mathbb{N}$ *the* $\mathbf{FI}^{\times \ell}$*-module* $\mathbf{H}^{\mathbb{Z}}_{2,n}(\underline{\Gamma}, \mathcal{Z})$ *is finitely generated in degree* $(n+6, \dots, n+6)$*.*

Proof of Theorem H. This follows from the special case of Theorem 4.25 where we have $Z_i = \varnothing$ for all $0 \leqslant i \leqslant \ell$. □

Proof of Theorem 4.25. Let $n \geqslant 2$, let $k_i \geqslant n+6$ for each $1 \leqslant i \leqslant \ell$, let $G = \mathbf{G}(k_1, \dots, k_\ell)$ and let $Z = Z_G$ be as defined above. We will prove that the homology of the configuration space of (G, Z) is generated by classes meeting only $n+6$ copies of each of the G_i. Here, meeting a copy of G_i means meeting the part which is not contained in G_0, i.e. $G_i - K_i$.

We inductively argue by removing some of the vertices in $V(G_0)$ and turning other such vertices into sinks. We will denote by (G', Z') a result of taking (G, Z) and applying any number of these two operations. Each vertex and edge of G' still belongs to a copy of one of the G_i. We will show that for each such pair (G', Z') the

second homology group $H_2(\mathrm{Conf}_n(G', Z'))$ is generated by classes meeting at most $n+6$ copies of each G_i. We will construct cycles where the moving particles meet only a bounded number of copies and argue that the non-moving particles only meet at most one copy of one G_i° each. Therefore, we will from now on assume that there are no isolated vertices or isolated edges without sinks, as the configuration spaces of these components have contractible connected components. Denote by G_i° the graph $G' \cap G_i$ where we removed all edges of K_i and all vertices which are not incident to an edge of $G_i - K_i$.

The induction is based on the number of essential non-sink vertices of $G_0 \cap G'$. For each $1 \leqslant i \leqslant \ell$ order the copies of G_i° in an arbitrary way, denoting the copies by G_{ij}° for $1 \leqslant j \leqslant k_i$. We may then talk about *the first* m *copies of* G_i° using this ordering. For each $m \leqslant n+6$ the intersection of G' with the union of G_0 and the first m copies of all G_i° will be called the m*-th base graph* $L_m \subset G'$. In particular, we have $L_0 = G_0 \cap G'$.

Induction start: The base case is that every essential vertex of L_0 is a sink. We will now describe a generating set for the second homology of configurations in the graph G' consisting of classes meeting only $n+3$ copies of each G_i°. We will construct this by describing a generating set consisting of classes represented by cycles where the moving particles meet at most five copies, and since we need at least two moving particles to form a 2-class this will imply that we in total meet at most $5 + n - 2 = n + 3$ copies. Recall that we forgot about the isolated edges and vertices because they always contain only fixed particles.

Look at the open cover of G' given by small contractible neighborhoods V_e of the edges $e \in E(L_0)$ and for each G_{ij}° the union of that graph and small open balls on the vertices in $V(L_0) \cap V(G_{ij}^\circ)$. We now describe the entries $E^\infty_{2,0}$, $E^\infty_{1,1}$ and $E^\infty_{0,2}$ of the corresponding Mayer-Vietoris spectral sequence for $\mathrm{Conf}_n(G', Z')$. Notice that the configuration space of the chosen neighborhood of each G_{ij}° can be shrunk to the configuration space of G_{ij}° because all vertices $V(L_0) \cap V(G_{ij}^\circ)$ are sinks, so we can pull the particles outside of G_{ij}° onto those sink vertices, see Figure 4.8.

The module $E^\infty_{0,2}$ admits a surjection from $E^1_{0,2}$, which is generated by

- 2-classes in some G_{ij}°,
- products of two 1-classes in some G_{ij}° and $G_{i'j'}^\circ$,
- products of a 1-class in some G_{ij}° and an H-class in L_0, and
- products of two H-classes in L_0,

tensored with zero-dimensional classes. The moving particles of these classes meet at most two copies of each G_i° because each H-class in the list above is contained

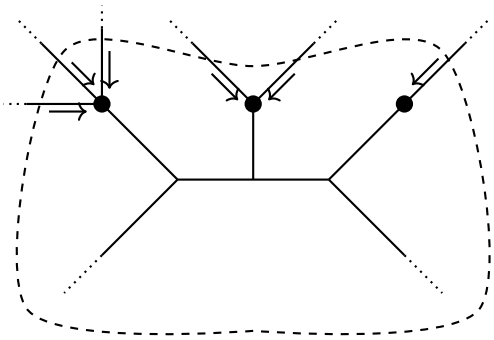

Figure 4.8: Pulling particles onto the sinks to show that the configuration space of a small neighborhood of G_{ij}° is homotopy equivalent to the configuration space of G_{ij}°.

in L_0, so the vertices involved in the H-class are sinks. There are at most $n-2$ non-moving particles, and they could possibly all be on different copies of the G_i°. Hence, we meet at most $2+n-2=n$ copies.

The module $E^\infty_{2,0}$: The module $E^1_{2,0}$ is given by direct sums of the tensor products of the zeroth homologies of configuration spaces of the open sets. Each connected component of G_{ij}° contains a sink vertex or an essential vertex, so it has connected configuration spaces. This does not change if we turn the whole connected component into a sink. Furthermore, H_0 does not see the loops inside G_{ij}°, which is why we now replace each G_{ij}° by a smaller graph.

We construct the following modified version M of G', which replaces each G_{ij}° by a smaller graph $\widehat{G}_{ij}^\circ$ that still connects the same vertices of $G_{ij}^\circ \cap L_0$ with each other as G_{ij}°: replace for each G_{ij}° each connected component intersecting L_0 in a set of vertices by the star graph given by the cone on those vertices. Each connected component of each G_{ij}° whose intersection with L_0 is empty is collapsed to a point (recall that each such component is assumed to contain an essential vertex or a sink vertex). It is straightforward to check that the module $E^2_{2,0}$ is the same as the module ${}^M E^2_{2,0}$ of the spectral sequence for $\mathrm{Conf}_n(M,M)$ with the analogous open cover given by small neighborhoods of edges in L_0 and small neighborhoods of the $\widehat{G}_{ij}^\circ$, the graph M was constructed in such a way that this is true. The configuration spaces of these open sets and their intersections have by construction contractible connected components, so all rows except for the zeroth row ${}^M E^1_{\bullet,0}$ are trivial. Therefore, ${}^M E^2_{2,0}$ is $H_2(\mathrm{Conf}_n(M,M)) = H_2(M^n)$, and thus generated by products of two particles moving along embedded circles. This implies that also $E^2_{2,0}$ is generated by all products of two particles moving along embedded circles in G', represented as linear

combinations of the particles being fixed in the intersections of open sets. We can choose a generating system consisting of such products where each circle meets only L_1 and at most one other copy of one of the G_i°. Mapping such a product via d_2 produces H-classes (one for each segment of the intersection of the two circles) and star classes (one for each isolated vertex in that intersection).

Now assume that we have a linear combination $\Sigma\alpha_i X_i$ of such elements in the kernel of

$$d_2 \colon E^2_{2,0} \to E^2_{0,1}.$$

By splitting this linear combination we can assume that each particle is for all X_i in the same path component of G'. Furthermore, each non-moving particle in the path component of L_0 can be assumed to be on one arbitrarily chosen sink vertex in L_0. Fix $i_0 \geqslant 0$ and $j_0 \geqslant 3$. The sum of all $\alpha_i X_i$ with X_i having two moving particles in $G^\circ_{i_0 j_0}$ has only classes inside L_1 because none of the other X_i produce classes in $G^\circ_{i_0 j_0}$: the intersection of all other pairs of circles intersects $G^\circ_{i_0 j_0}$ in a subset of the set of vertices $L_0 \cap G^\circ_{i_0 j_0}$, and since all vertices of L_0 are sinks there are no star classes at these vertices. Now look at the same sum of elements with the moving particles moved to the second copy $G^\circ_{i_0 2}$ instead of $G^\circ_{i_0 j_0}$, then this produces the same sum of classes in L_1. Therefore, their difference is an element of the kernel, and subtracting such elements we can assume that all X_i having both moving particles in the same copy of any G_i° only meet the second base graph L_2. The moving particles of the kernel classes we subtracted meet at most three copies of each G_i°, and together with the $n-2$ non-moving particles these classes meet not more than $3+n-2 = n+1$ copies.

Fixing two distinct copies $G^\circ_{i_0 j_0} \neq G^\circ_{i_1 j_1}$ we now look at an X_i where two particles move in those copies, one particle per copy. Under d_2, this again only produces classes in L_1 because the intersection of the two circles is contained in L_1. The same product moved to the second (and for $i_0 = i_1$ also third) copies produces the same sum of classes in L_1, so their difference is an element of the kernel. The moving particles meet in this difference at most five copies of each G_i°, and together with the non-moving particles the difference meets at most $5+n-2 = n+3$ copies.

By subtracting such kernel elements we can arrange that the moving particles of each X_i meet only L_3. By grouping the X_i according to the path components in which the different particles are we can write $\Sigma\alpha_i X_i$ as a sum of kernel elements meeting at most $3+n-2 = n+1$ copies of each G_i°.

This shows that $E^\infty_{2,0} \cong E^3_{2,0}$ is generated by classes meeting at most $n+3$ copies of each G_i°. By their description, it is straightforward to check that the lifting required to interpret these as homology classes in $H_2(\mathrm{Conf}_n(G', Z'))$ does not increase the number of copies met, see Section 2.1.1.

The module $E^\infty_{1,1}$ is a quotient of $E^2_{1,1}$, so we only have to give generators of that

latter module. Let M be defined as above, and define M_{ij} in the same way, except that G°_{ij} is left untouched. By the same arguments as in the previous case we get

$$E^2_{1,1} \cong \bigoplus_{e \in E(G_0)} {}^{M_e}E^2_{1,1} \oplus \bigoplus_{i,j} {}^{M_{ij}}E^2_{1,1},$$

where ${}^{M_e}E^*_{\bullet,\bullet}$ is the spectral sequence for $\mathrm{Conf}_n(M, M - \mathring{e})$ and ${}^{M_{ij}}E^*_{\bullet,\bullet}$ is the spectral sequence for $\mathrm{Conf}_n(M_{ij}, (M_{ij} - G^\circ_{ij}) \cup (Z_{G'} \cap G^\circ_{ij}))$. The open covers are again given by neighborhoods of edges and neighborhoods of the G°_{ij} and $\widehat{G}^\circ_{ij}$. We now give generators of the modules on the right-hand side. For easier description, we will in the following say that a particle meets some G°_{ij} even if we replaced that graph and the particle meets $\widehat{G}^\circ_{ij}$.

To calculate the homology of $\mathrm{Conf}_n(M, M - \mathring{e})$, let F be a forest consisting of maximal trees in the connected components of $M - \mathring{e}$ and collapse all components of the resulting spanning forest of M to point shaped sinks. Via this collapse, $\mathrm{Conf}_n(M, M - \mathring{e})$ is homotopy equivalent to the configuration space of the disjoint union of isolated sinks, wedges of copies of (S^1, S^1), and either

- the wedge of (S^1, pt) and copies of (S^1, S^1) or
- the wedge of $([0,1], \{0,1\})$ and copies of (S^1, S^1) at $0, 1 \in [0,1]$,

see Figure 4.3 and Figure 4.9. All wedges above are taken at the points $1 \in S^1$ and $0, 1 \in [0,1]$. By Proposition 3.12 the homology of these configuration spaces is generated by products of classes of configurations in the individual wedge summands. In $\mathrm{Conf}_n(M, M - \mathring{e})$, this gives products of H-classes and circle classes such that at most one of them meets the edge e.

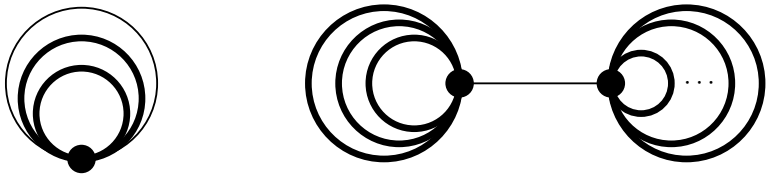

Figure 4.9: Collapsing a forest of sink edges.

The classes involving an H-class at e are realized in ${}^{M_e}E^2_{1,1}$, and to generate all such classes it suffices to take those where the moving particles meet only 2 copies of each G°_i: the H-part is in L_0 and we meet two copies for the circle part (as above, circle classes can be assumed to meet only the first base copies and one additional copy). All other classes are realized in the bottom row, so the remaining part of ${}^{M_e}E^2_{1,1}$ is hit by the boundary map

$$d_2 \colon {}^{M_e}E^2_{3,0} \to {}^{M_e}E^2_{1,1}.$$

The module ${}^{M_e}E^2_{3,0}$ is isomorphic to $H_3(M^n)$ by comparing it to the corresponding module of the spectral sequence for $\mathrm{Conf}_n(M, M)$, and is therefore generated by products of three particles moving along embedded circles. We can again assume that each of those embedded circles meets apart from the first base copies at most one G°_{ij}, so the moving particles of each such product meet at most 4 copies of each G°_i.

Therefore, the direct summand of $E^2_{1,1}$ corresponding to ${}^{M_e}E^2_{1,1}$ is generated by elements meeting at most $4+n-2 = n+2$ copies of each G°_i. Again, the interpretation as homology classes in $H_2(\mathrm{Conf}_n(G', Z'))$ does not increase the number of copies met, see Section 2.1.1.

To compute ${}^{M_{i_0j_0}}E^2_{1,1}$, collapse a maximal forest in $M_{i_0j_0} - (G_{i_0j_0} - K_{i_0})$, yielding a disjoint union of isolated vertices, wedges of copies of (S^1, S^1), and the wedge sum of a quotient $\overline{G^\circ_{i_0j_0}}$ of $G^\circ_{i_0j_0}$ and copies of (S^1, S^1), see Figure 4.10. Here, the circle S^1 is always wedged at $1 \in S^1$. We will call the union of L_0 and all $\widehat{G}^\circ_{ij}$ such that $j \leqslant m$ and $(i, j) \neq (i_0, j_0)$ the m-th base graph $\overline{L}_m$ of $M_{i_0j_0}$. We choose the maximal forest above such that restricted to each connected component of the first base graph it is a spanning tree. As above, the second homology of the configuration space of this

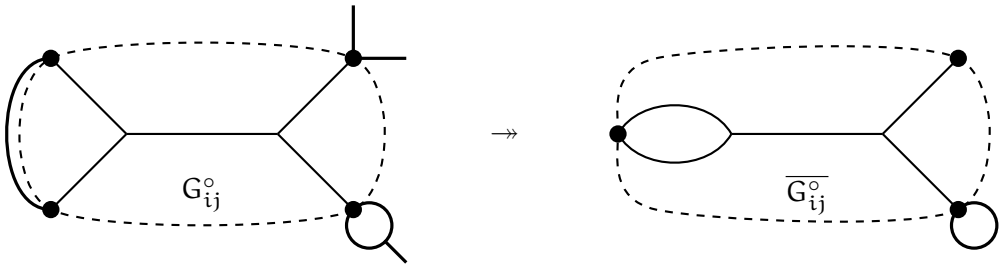

Figure 4.10: Collapsing a forest of sink edges.

quotient of $(M_{ij}, (M_{ij} - G^\circ_{ij}) \cup (Z_{G'} \cap G^\circ_{ij}))$ is generated by

- 2-classes in $\overline{G^\circ_{i_0j_0}}$,
- products of 1-classes in $\overline{G^\circ_{i_0j_0}}$ and circle classes in the copies of S^1, and
- products of two circle classes in the copies of S^1,

with the remaining particles fixed somewhere in the graph. The classes in $\overline{G^\circ_{i_0j_0}}$ are represented in $M_{i_0j_0}$ by classes whose moving particles meet only $G_{i_0j_0}$ and the base graph $\overline{L}_1$ by the choice of maximal forest. The moving particles of the circle classes meet at most one copy of one G°_i outside of the base graph $\overline{L}_1$. Therefore, the moving particles of all classes above meet at most two copies outside of the base

graph $\overline{L}_1$. We will now describe how these classes are distributed onto the different modules of the infinity page ${}^{M_{i_0j_0}}E^{\infty}_{\bullet,\bullet}$.

The products of the third kind are elements in ${}^{M_{i_0j_0}}E^{\infty}_{2,0}$. We now look at products of the second kind. The 1-class in $\overline{G^{\circ}_{i_0j_0}}$ can by Theorem D assumed to be a star class, an H-class or a circle class. All star classes and H-classes in $\overline{G^{\circ}_{i_0j_0}}$ come from star classes and H-classes in $G^{\circ}_{i_0j_0}$. Therefore, such products are represented as elements in ${}^{M_{i_0j_0}}E^{\infty}_{1,1}$. For circle classes coming from circle classes in $G^{\circ}_{i_0j_0}$ the same is true. All other circle classes move inside $G^{\circ}_{i_0j_0}$ and $\overline{L}_1$ by construction of the forest. The product of such a class with a circle class in one of the S^1's is represented in ${}^{M_{i_0j_0}}E^{\infty}_{2,0}$.

Therefore, ${}^{M_{i_0j_0}}E^{\infty}_{1,1}$ is generated by the products of circle classes in the copies of S^1 with 1-classes coming from 1-classes in $G^{\circ}_{i_0j_0}$ as described above, and possibly elements representing the first kind of classes listed above. The moving particles of all these elements meet at most three copies of each G°_i. Since the image of d_2 landing at position $(1,1)$ is generated by elements whose moving particles meet the image of at most four copies of each G°_i as above, this shows that the direct summand of $E^2_{1,1}$ corresponding to ${}^{M_{i_0j_0}}E^2_{1,1}$ is generated by classes meeting at most $4+n-2 = n+2$ copies of each G°_i.

Therefore, $E^{\infty}_{1,1}$ is generated by classes meeting at most $n+2$ copies of each G°_i, and the same is true for the corresponding classes in the second homology of $\mathrm{Conf}_n(G', Z')$.

Induction step: Assume that we have an essential vertex v in L_0 which is not a sink vertex. Look at the open cover given by a small contractible open ball V_v around v and $V_r := G' - \{v\}$.

We describe generators for the second co-diagonal of the corresponding Mayer-Vietoris spectral sequence for the induced cover of $\mathrm{Conf}_n(G', Z')$.

The module $E^{\infty}_{0,2}$ is generated by 2-classes in V_r and products of 1-classes in V_v and V_r. By induction, the 2-classes in V_r can be written as sums of elements meeting only $n+6$ copies of each of the G°_i. The products of 1-classes consisting of k_1 and k_2 particles with $k_1+k_2+k_3 = n$ can be chosen to each meet at most k_j+3 copies of each of the G°_i by Theorem F. The remaining k_3 particles are part of zero-classes and therefore meet at most k_3 copies in total, so such a product meets at most $k_1+3+k_2+3+k_3 = n+6$ copies of each G°_i.

The module $E^{\infty}_{2,0}$: The module $E^2_{2,0}$ is the same as the module ${}^{v}E^2_{2,0}$ of the spectral sequence ${}^{v}E^*_{\bullet,\bullet}$ for the quotient where all edges disjoint from V_v and incident to an essential vertex or a sink vertex are collapsed to sink vertices, see Figure 4.11. This is the disjoint union of

- isolated sink vertices,

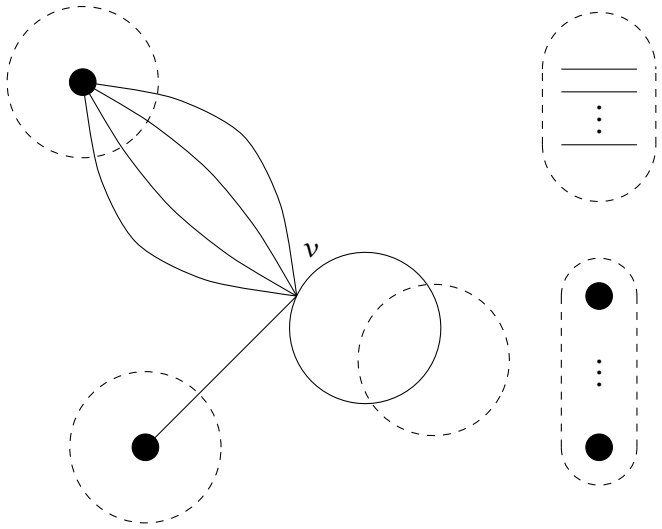

Figure 4.11: Collapsing edges contained in V_r gives a disjoint union of vertices, edges and one wedge of pointed banana graphs. The dashed lines indicate the image of V_r under the collapse map.

- isolated edges *not incident to sink vertices*, and
- the wedge of pointed banana graphs (possibly with only one or two edges) where all essential vertices except the central vertex v are sinks. Banana graphs with one or two edges can either have one or zero sinks.

Therefore, by Proposition 1.20 the configuration space of this quotient is 1-dimensional. Under this identification the kernel of

$$d_2 \colon E^2_{2,0} \to E^2_{0,1}$$

is contained in the kernel of

$$d_2 \colon {}^{v}E^2_{2,0} \to {}^{v}E^2_{0,1}.$$

But this latter kernel has to be trivial by the dimension estimate from above, which means that $E^\infty_{2,0} = E^3_{2,0} = 0$.

The module $\mathbf{E^\infty_{1,1}}$ is a quotient of $E^2_{1,1}$, which is given by the direct sum of the modules ${}^{v}E^2_{1,1}$ and ${}^{\tau}E^2_{1,1}$, where ${}^{\tau}E^*_{\bullet,\bullet}$ is the spectral sequence for $(G', Z' \sqcup \{v\})$ and ${}^{v}E^*_{\bullet,\bullet}$ is defined as in the previous case.

As described above, the spectral sequence ${}^{v}E^*_{\bullet,\bullet}$ converges to the homology of a configuration space with homological dimension 1, so ${}^{v}E^2_{1,1}$ is given by the image of the map

$$d_2 \colon {}^{v}E^2_{3,0} \to {}^{v}E^2_{1,1}.$$

To compute ${}^{\nu}E^2_{3,0}$ we can turn the only non-sink vertex ν into a sink as well. The spectral sequence for that graph is only non-trivial in the zeroth row, so ${}^{\nu}E^2_{3,0}$ is the third homology of the graph described in the previous case with ν turned into a sink. This is (after forgetting the isolated intervals and vertices) the wedge sum of pointed banana graphs (possibly with only one or two edges) with all essential vertices sinks, so by Proposition 3.12 the third homology is generated by products of classes of configurations in the individual banana graphs. By Proposition 4.14, the 3-classes in one such banana graph can be written as sums of classes where the moving particles meet at most 8 edges. By Proposition 4.13, the second homology of configurations in one banana graph is generated by classes meeting at most 4 edges. The first homology is generated by H-classes and S^1 classes by Theorem D, so it can be generated by classes where the moving particles meet at most 2 edges. Taking the at most $n-3$ non-moving particles into account, ${}^{\nu}E^2_{3,0}$ and therefore ${}^{\nu}E^2_{1,1}$ can be generated by classes meeting at most $8+n-3 = n+5$ edges.

To interpret those classes as elements in the configuration space of (G', Z'), we have to lift a class in $E^2_{1,1}$ by mapping a representative on the zeroth page via d_1 and lifting the result along d_0, see Section 2.1.1. The corresponding lifting in ${}^{\nu}E^*_{\bullet,\bullet}$ exchanges particles over one of the newly added sinks, so in $E^*_{\bullet,\bullet}$ we have to move and reorder fixed particles along the connected components that we collapsed to sinks. It is straightforward to check that this can be done by meeting only the $n+5$ copies of each G_i° from above and *one* additional copy for each i. Therefore, these classes meet at most $n+6$ copies of each G_i°.

The spectral sequence ${}^{r}E^*_{\bullet,\bullet}$ converges to the homology of configurations in $(G', Z' \sqcup \{\nu\})$. By induction, the second homology of that space is generated by classes meeting only $n+6$ copies of each G_i°. We now want to show that the same is true for ${}^{r}E^2_{1,1}$ by showing that this holds for ${}^{r}E^\infty_{0,2}$ and ${}^{r}E^\infty_{2,0}$.

It is true by induction for ${}^{r}E^2_{0,2}$ and therefore ${}^{r}E^\infty_{0,2}$ since the former only consists of 2-classes of particles in V_r, so it remains to show this for ${}^{r}E^2_{2,0}$.

Collapsing the edges disjoint from V_ν as above, we see that ${}^{r}E^2_{2,0}$ is the second homology of the configuration space of a disjoint union of isolated vertices, isolated edges without sinks and the wedge of pointed banana graphs (with possibly only one or two edges) with *all* essential vertices sinks. The isolated vertices and edges can only be involved in zero-classes, so we can focus on the wedge of banana graphs. By Proposition 3.12 and Proposition 4.14, the module ${}^{r}E^2_{2,0}$ is thus generated by products of classes in the individual wedge summands. We will now show that the kernel of

$$d_2 \colon {}^{r}E^2_{2,0} \to {}^{r}E^2_{0,1}$$

can be generated by classes meeting at most $n+6$ copies of each G_i°.

A product of two 1-classes (either circle classes or H-classes) in different wedge

summands maps to zero under d_2 and therefore represents an element of the ∞-page. Interpreting it as a class in $\mathrm{Conf}_n(G', Z' \sqcup \{v\})$ we end up with a class where the moving particles meet at most five copies of each G_i°, so considering the non-moving particles the class meets at most $5 + n - 2 = n + 3$ copies.

It remains to understand the 2-classes in one fixed banana graph. We can assume that no fixed particle is in any other banana graph by moving them onto the sink v. From the proof of Proposition 4.13 we get a description of generators of $H_2(\mathrm{Conf}_n(B_k))$ as generators for the modules of the infinity page of the spectral sequence used there.

If the banana graph has at most four edges, then by $k_i \geqslant n + 6 > 4$ we see that all of these edges must be edges of L_0. Assume that we have an element of the infinity page with $m \geqslant 2$ particles in the banana graph. We want to interpret it as an element of $H_2(\mathrm{Conf}_n(G', Z' \sqcup \{v\}))$, for which we have to choose lifts. Since we have an element of the infinity page, we know that these lifts exist. Choose m copies of each of the G_i°, one copy for each particle in the banana graph. We can modify each of these lifts in the following way: whenever a particle moves along some edge of some G_{ij}°, we replace this movement by the movement in the copy of G_i° we just chose for this specific particle. It is straightforward to check that this still bounds the same cycle, and additionally the final result only meets at most m copies of each G_i°. Together with the $n - m$ particles in the other connected components of $(G', Z' \sqcup \{v\})$, this gives classes meeting at most n copies of each G_i°.

If the banana graph has five or more edges, then only the module at position $(2, 0)$ of the spectral sequence in Proposition 4.13 is non-trivial, and it is generated by tori and surfaces of genus 2. These come from products of two particles moving along embedded circles, and in the genus 2 surface case we glued two of those together where each product has one edge used by both particles, see the proof of Proposition 4.13. In particular, for each of those classes the remaining $n - 2$ particles sit on the vertex v or on the isolated vertices and edges. For two of these classes to land under d_2 in the same direct summands the fixed particles have to define the same zero-class in the graph, so we only need to consider kernel elements given by linear combinations of such classes where the fixed particles all determine the same zero-class. Therefore, by projecting to pairs of particles we can assume that we actually have $n = 2$, the general case follows by adding the fixed particles again.

Let $\Sigma \alpha_i X_i$ be a linear combination of elements of these two types in the kernel of d_2. We will now see that this linear combination can be decomposed into linear combinations in the kernel of d_2 such that each of those linear combinations uses edges of at most seven copies of each G_i°. Fixing two distinct edges e_1 and e_2 of the banana graph contained in L_2 and adding tori to exchange free edges of the circles we can arrange that each X_i meets e_1, e_2 and at most two more edges, see Remark 4.16. Now assume we have i_0 such that X_{i_0} meets two edges $e_1^{i_0}, e_2^{i_0}$ not

contained in L_4. Define for each particle $m \in \{1,2\}$ a continuous map $\xi_m \colon G' \to G'$ sending every G°_{ij} via the identity to $G^\circ_{i(m+2)}$. Now define the following continuous self-map $\mathrm{Conf}_2(G', Z' \sqcup \{v\}) \to \mathrm{Conf}_2(G', Z' \sqcup \{v\})$: if particle m is in a copy of some G°_i outside of L_2 and not containing one of the two edges $e_1^{i_0}, e_2^{i_0}$, then change its position by applying ξ_m. This induces a self-map of the Mayer-Vietoris spectral sequence because ξ_1 and ξ_2 map V_v, V_r and $V_v \cap V_r$ into themselves, so the image of our element in the kernel of d_2 is again a kernel element. By construction, all X_i meeting only L_2 and the copies containing $e_1^{i_0}$ or $e_2^{i_0}$ are mapped by the identity. The images of the remaining X_i meet at most one copy of one G°_i outside of L_4. Therefore, we can replace all elements meeting both copies containing $e_1^{i_0}$ and $e_2^{i_0}$ by elements meeting only one copy outside of L_4, and in this replacement we added a kernel element meeting at most six copies of each G°_i. Repeating this argument, we can assume that all X_i meet at most one edge outside of L_4.

By an analogous argument we can subtract kernel elements meeting at most seven copies of each G°_i to arrange that all X_i only meet edges in L_6, thus meeting at most six copies of each G°_i. To interpret such a class with two moving particles as a class in $H_2(\mathrm{Conf}_n(G', Z' \sqcup \{v\}))$ we have to choose lifts in V_r again. The classes meeting at most seven copies already come with a choice of copies of each G°_i for each of the two particles such that this copy is not met by the other particle. The other classes meet at most six edges and we choose for each of the two particles one additional copy of each G°_i, so that we have at most eight copies. We can then use the same technique as for the case of banana graphs with at most four edges to see that we can modify the lifts to not meet more than these eight copies.

The fixed particles each meet at most one copy of one G°_i each, so we constructed generators meeting at most $8 + n - 2 = n + 6$ copies of each G°_i.

This shows that also ${}^rE^\infty_{2,0} = {}^rE^3_{2,0}$ is generated by classes meeting at most $n + 6$ copies of each G°_i, so the same is true for ${}^rE^3_{1,1} = {}^rE^\infty_{1,1}$. By a similar argument as for ${}^vE^\infty_{1,1}$, the image of

$$d_2 \colon {}^rE^2_{3,0} \to {}^rE^2_{1,1}$$

is generated by classes meeting at most $8 + n - 3 = n + 5$ copies of each G°_i, and therefore ${}^rE^2_{1,1}$ is generated by classes meeting at most $n + 6$ copies of each G°_i. To interpret these classes in the configuration space of (G', Z'), we have to lift them by choosing paths of particles inside V_v. We only have to connect different zero cycles in V_v, and to do this we need at least three edges. These edges can be chosen to be in copies of G°_i we already met, or if we did not meet three copies of any G°_i we choose three edges at random and still meet at most $n + 6$ copies of each G°_i.

Hence, the module $E^2_{1,1}$ is generated by classes meeting at most $n + 6$ copies of each G°_i, concluding the induction step. □

4.5 Stability for increasing number of particles

In this section we will prove that for highly connected graphs the map forgetting the last particle induces in first cohomology a representation stable sequence of representations of the symmetric group.

Let G be a graph, let S, T be finite sets and let $\iota\colon S \hookrightarrow T$ be an injection. Then we get an induced map

$$\mathrm{Conf}_T(G) \to \mathrm{Conf}_S(G)$$

by precomposition. This gives $\mathrm{Conf}_\bullet(G)$ the structure of an $\mathbf{FI}^{\mathrm{op}}$-space. The cohomology functor $H^1(-;A)$ for an abelian group A turns this into an **FI**-module

$$H^1(\mathrm{Conf}_\bullet(G);A)$$

over A. Recall that a graph is k-vertex connected if every pair of vertices (v,w) can be connected by k distinct paths intersecting only in their endpoints.

Theorem I. *Let G be a finite 3-vertex connected graph with at least four essential vertices and without self-loops. Let A be an abelian group such that $H_1(\mathrm{Conf}_2(G);A)$ is torsion-free. Then $H_1(\mathrm{Conf}_n(G);A)$ is torsion-free for all n and the* **FI***-module $H^1(\mathrm{Conf}_\bullet(G);A)$ is finitely generated in degree 2. In particular, the sequence $n \mapsto H^1(\mathrm{Conf}_n(G);\mathbb{Q})$ induced by forgetting the last particle is representation stable and its dimension is eventually polynomial in n.*

High vertex connectivity has the following implication: let G be k-vertex connected for $k \geqslant 1$ and let $v, v', w_1, \dots, w_{k-1} \in V(G)$ be all distinct. Then there exists an edge path between v and v' not meeting $w_1, \dots, w_{k-1}$: there exist at least k paths between v and v' only intersecting in their endpoints, and at most $k-1$ of them can meet one or more of the w_i.

For the proof of Theorem I, we need the following results.

Proposition 4.26. *Let G be as in Theorem I, then the first homology $H_1(\mathrm{Conf}_n(G);A)$ is generated by basic classes with at most two moving particles.*

Proof. By Theorem C, the first homology group $H_1(\mathrm{Conf}_n(G);A)$ can be generated by S^1-, H- and star classes. Since we have $G \neq S^1$, we saw in the proof of that theorem that we in fact only need S^1-classes involving *one* moving particle. By Proposition 1.24 it suffices to take H-classes with two moving particles, so it remains to show that we can write every star class as a linear combination of classes with two moving particles.

Let Z be a cellular 1-cycle in the combinatorial model of $\mathrm{Conf}_S(G)$ representing a non-trivial star class at v whose image is an embedded circle in the 1-skeleton of $\mathrm{Conf}_S(G)$. Let $0 \leqslant \ell \leqslant n-3$ be the number of non-moving particles of Z, then we

will show that we can write Z as a linear combination of classes where at least $\ell + 1$ particles are fixed. Fix one edge e_0 which is *not* contained in the star of v and move all fixed particles onto that edge. This edge exists because of connectivity and the fact that we have at least three vertices.

Let $s \in S$ be any particle and look at a connected segment γ of Z such that the 0-cube at the beginning of the segment denoted by γ_0 has the particle s on v, the next 1-cube moves it onto an edge e and the particle s is back on v for the first time at the end of the segment, which we denote by γ_1. If F is the smallest closed interval in e_0 containing all fixed particles, then by the connectivity assumption the graph $G - \{v\} - F$ is connected and contains at least one essential vertex, so it has connected configuration spaces. Therefore, we can choose a path γ' in $\mathrm{Conf}_S(G)$ connecting the configurations γ_0 and γ_1 such that s is fixed on v and the ℓ particles which are fixed for Z are also fixed for γ'. The path γ followed by the inverse of γ' is homologous to a cycle where s is fixed on the edge e: s is already fixed on e for γ, except for the very first and last 1-cubes. The path γ' extended by those two cubes is homologous relative endpoints to the same path γ' with s moved from v onto the edge e. This gives a class with $\ell + 1$ fixed particles.

The path Z' given by Z with γ replaced by γ' moves s strictly less, so repeating this process we eventually reach a cycle that does not move s at all and therefore has at least $\ell + 1$ fixed particles as claimed. Notice that Z' may not be a circle anymore, but we can always write it as a sum of embedded circles in the combinatorial model without increasing the number of movements of s. Furthermore, the particle s never leaves the star of v, which allows the repetition of the argument above.

Each of those classes with at least $\ell + 1$ fixed particles can now be written as linear combinations of basic classes with at least $\ell + 1$ fixed particles: For a class with $m \geqslant \ell + 1$ fixed particles move all those fixed particles onto edges, cut the graph at those m positions and project to the configuration space with only the non-fixed particles. This gives a new class in the cut graph, and we can write it as a sum of basic classes. Then, we include this graph into G again and add the fixed particles to their original positions. All circle classes involve one moving particle again and every H-class can be written as a sum of H-classes with two moving particles. For star classes, repeat the argument above until there are at least $n - 2$ fixed particles. □

Proposition 4.27. *Let G be as in Theorem I. Then there exists a basis of $H_1(G)$ consisting of embedded circles such that none of the circles meets all vertices.*

Proof. To see this, it is sufficient to construct a spanning tree with a vertex of valence at least three, the corresponding basis of the homology then has the required property. Choose four distinct vertices v_1, v_2, v, v' such that v_1 and v_2 are connected by an edge e_1. We can find a path from v_1 to v not meeting v_2. Let e_2 be the first edge in that path and v_3 the other vertex incident to e_2. Now choose a path connecting v_1

and v' without meeting v_2 and v_3, and let e_3 be the first edge in that path. Then $e_1 \cup e_2 \cup e_3$ is a star graph with three edges, and by extending it to a spanning tree we get a spanning tree with a vertex of valence at least three. □

Proposition 4.28. *Let* G *be as in Theorem I. Let* X *be a cycle representing a circle class* $[X] \in H_1(\mathrm{Conf}_n(G))$ *with one particle* s *moving along an embedded circle in* G *and* $n-1$ *fixed particles. For each edge* e *denote by* F_e *the set of fixed particles on that edge. If the complement of the circle contains apart from* e *at least two more edges, then* $[X] = [\sigma X]$ *for each* $\sigma \in \Sigma_{F_e}$.

Proof. Choose two distinct edges e_1, e_2 in the complement of the embedded circle which are also distinct from e. Let v and w be the two vertices incident to e. Choose paths γ_1 and γ_2 from v to e_1 and e_2, respectively, ending in vertices we denote by v_1 and v_2. By connectivity, we can choose these paths such that the complement of each γ_i contains w and at least one vertex of each e_1 and e_2 (if γ_i meets both vertices of e_i, then we can restrict to a subpath that does not).

Every reordering of fixed particles on e can now be done (ignoring the moving particle s for now) by moving individual fixed particles back and forth along γ_1 and γ_2 and the edges e, e_1 and e_2. Since we are only changing the order of fixed particles *inside* the edge e, each particle moves the same number of times γ_i as the inverse of γ_i. The corresponding paths in the configuration space moving a particle along γ_i and the inverse of γ_i only differ in the positions of the other fixed particles.

It is straightforward to check that each of these changes of position creates

- at each edge in the intersection of γ_i and the embedded circle an H-class of s and the fixed particle whose position we change, and
- at each isolated vertex in this intersection a star class of those two particles, see Figure 4.12.

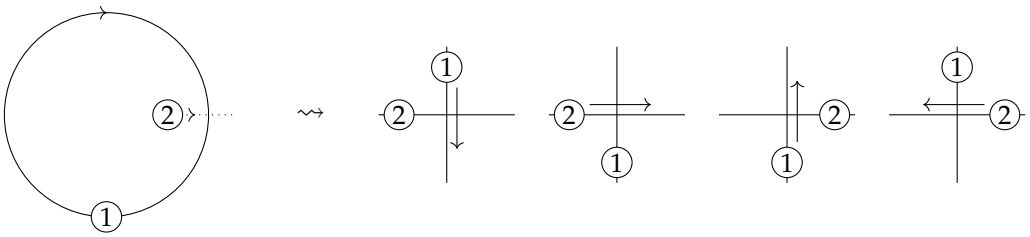

Figure 4.12: Moving a fixed particle along a path intersecting the circle class creates a star class at that vertex. Moving it back creates the additive inverse of this class.

If we ignore the other fixed particles, then the classes arising by moving back and forth along one of the γ_i cancel each other out. For star classes, the position of the

other fixed particles does not matter because G without any vertex has connected configuration spaces, so those homology classes also cancel when we take the fixed particles into account. By construction, none of the H-classes meet both vertices of any of the edges e, e_1 or e_2. Therefore, we can find paths from the edges e_1 and e_2 to e which do not meet the vertices involved in the H-class. Along these paths we can move all fixed particles onto e and change their order by using one of the vertices v and w not involved in the H-class without changing the homology class. Hence, we can arrange that the position of all fixed particles of these H-classes is always the same, so the H-classes cancel as well when taking the fixed particles into account. □

Proposition 4.29. *Let* G *be as in Theorem I. Then each* H*-class in* $H_1(\mathrm{Conf}_n(G))$ *can be written as a linear combination of star classes with two moving particles.*

Proof. By Proposition 1.24, it suffices to show this for H-classes with two moving particles. Let such an H-class on an edge e incident to vertices v and w be given, and let e_f be an edge which is neither incident to v nor to w (this exists by connectivity and the fact that we have at least four vertices). We will first show that we can move all fixed particles onto the edge e_f by only adding star classes with two moving particles. By connectivity, we can move all fixed particles which are not on an edge incident to both v and w onto e_f without changing the homology class by moving them in the complement of a small neighborhood of e. We now describe what happens when we move a fixed particle s on an edge e_1 incident to v and w onto the edge e_f. Let u be one of the vertices of e_f, then there exists an edge path from v to u which does not meet w. Moving s onto the edge e_u given by the first edge in that path produces a star class at v of the following kind:

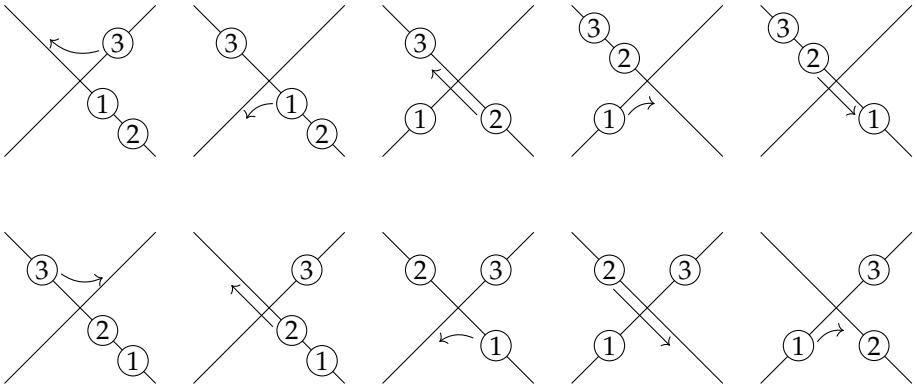

The central vertex in these pictures is v, the top left edge is e_u, the top right edge is e_1, the bottom right edge is e, and $s = 3$. Notice that this picture shows the special case where the valence of v is four and the H-class involves the two edges on the left, but by adding star classes with two moving particles we can always arrange this situation (forgetting about the remaining edges at v). Once s is moved to e_u, we can move it to e_f without changing the H-class. We will now show that the star class

above can be represented by a star class with two moving particles, which proves that up to adding such star classes every H-class can be assumed to have all particles on e_f.

By connectivity, there exists a path from v to w whose first edge is e_u, which meets v and w only at its endpoints and which does not meet the interior of e_f. Moving s along this path from e_u to e_1 we see that the star class can be written as the sum of a star class with s fixed on e_1 and two classes moving s along a circle with the two particles of the H-class fixed on e in the two different orderings:

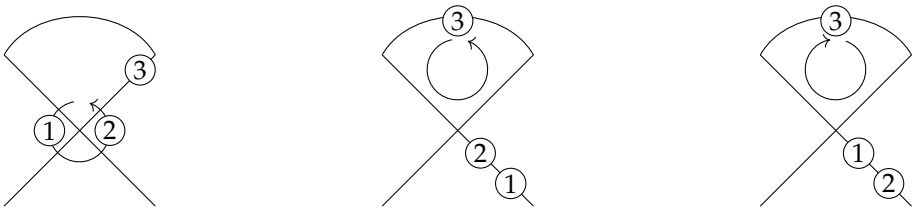

The first of these classes is a star class with two moving particles, so we can forget about it. Since we have at least four vertices of valence at least three and therefore at least three edges in the complement of any embedded circle, we see by Proposition 4.28 that the latter two classes are the same with different signs, so they add to zero.

Therefore, we can assume that the H-class has all fixed particles on e_f. We now show that each such H-class can be written as a linear combination of star classes with two moving particles.

There exist by connectivity two disjoint paths connecting v and w without meeting the interior of e. This gives an embedded theta graph (the banana graph with three edges) with trivalent vertices at v and w. If e_f is contained in this theta graph then choose a different edge e'_f whose interior is disjoint from the theta graph and which is not incident to both v and w. This exists because we have at least four vertices of valence at least three. Connecting e_f and e'_f via a path not meeting v and w we can move all fixed particles onto e'_f without changing the homology class. Therefore, we can forget about the fixed particles and just show that an H-class of two particles in the theta graph can be written as a linear combination of star classes. Adding star classes we can arrange that one particle only uses the first two edges of the theta graph and the other particle only uses the last two edges. For each movement of one particle along e, move the fixed particle into the middle of the edge of the theta graph it is on. It is straightforward to see that this gives four circle classes which add to zero. Therefore, we reached the trivial class.

This shows that the H-class in G can be written as a sum of star classes with two moving particles. □

Proof of Theorem I. By Proposition 4.5, the homology is generated by basic classes

with at most two moving particles. We can further assume that the circle classes only have one moving particle, and by Proposition 4.29 we do not need H-classes.

When we say that the first homology is generated by these types of classes, we always have to add the particles not involved in the 1-classes onto the graph in all possible ways: an embedding of one of those graphs K into G determines maps

$$H_1(\mathrm{Conf}_S(K); A) \otimes H_0(\mathrm{Conf}_{\underline{n}-S}(G-K); A) \to H_1(\mathrm{Conf}_n(G); A). \tag{4.1}$$

All these maps for $|S| \leqslant 2$ together generate the first homology. We will now argue that by the connectivity assumption it suffices to add the fixed particles in one single way instead.

Assume we are given a **circle class** moving one particle along an embedded circle and keeping the remaining $n-1$ particles fixed. We can move one of the fixed particles from one position to any other position in the complement of the circle via a path γ because G is connected. As described in the proof of Proposition 4.28, this creates H-classes and star classes in the intersections of γ and the embedded circle, see Figure 4.12. Therefore, we can choose for each embedding $S^1 \hookrightarrow G$ an arbitrary way of putting the remaining $n-1$ particles into the complement of the circle. Since the H-classes can again be written as linear combinations of star classes, these circle classes together with the star classes will still generate the whole homology.

For **star classes** it is clear that all ways of adding fixed particles give rise to the same homology class by the fact that $G - \{v\}$ has connected configuration spaces for all vertices v.

Each $H_1(\mathrm{Conf}_{\{s_1,s_2\}}(G); A)$ splits into modules of circle classes and star classes. More precisely, there is a map

$$H_1(\mathrm{Conf}_{\{s_1\}}(G); A) \oplus H_1(\mathrm{Conf}_{\{s_2\}}(G); A) \cong H_1(G; A)^{\oplus 2} \to H_1(\mathrm{Conf}_{\{s_1,s_2\}}(G); A)$$

given as follows: using Proposition 4.27 we choose a basis of $H_1(G; A)$ consisting of embedded circles such that none of them meet all vertices of G, and for each of those circles we choose an essential vertex in its complement. We now send each element of this basis of $H_1(G; A)$ to the same element with the other particle added to the chosen vertex. The composition of this map with the direct sum of the projection maps

$$\pi_{s_1} \oplus \pi_{s_2} \colon H_1(\mathrm{Conf}_{\{s_1,s_2\}}(G); A) \to H_1(\mathrm{Conf}_{\{s_1\}}(G); A) \oplus H_1(\mathrm{Conf}_{\{s_2\}}(G); A)$$

is the identity, so we have

$$H_1(\mathrm{Conf}_{\{s_1,s_2\}}(G); A) \cong H_1(\mathrm{Conf}_{\{s_1\}}(G); A) \oplus H_1(\mathrm{Conf}_{\{s_2\}}(G); A) \oplus H[2]_{\{s_1,s_2\}}$$

for some module $H[2]_{\{s_1,s_2\}}$. By the discussion above, this module is generated by star classes and maps to zero under the projections to a single particle.

Fixing a finite set S we now construct a map

$$\bigoplus_{s_1\in S} H_1(\mathrm{Conf}_{\{s_1\}}(G);A)\oplus \bigoplus_{\{s_1,s_2\}\subset S} H[2]_{\{s_1,s_2\}} \to H_1(\mathrm{Conf}_S(G);A) \tag{4.2}$$

as follows: choose bases of all $H_1(\mathrm{Conf}_{\{s_1\}}(G);A)$ by pulling back the basis we chose above for $H_1(G;A)$ under the canonical isomorphisms. Each circle class in these bases is mapped by adding the fixed particles onto an edge in the star of the vertex we chose above. For each summand $H[2]_{\{s_1,s_2\}}$ choose a basis consisting of linear combinations of star classes. Fix an arbitrary edge e_f of G and map each such basis element by adding the fixed particles onto the interior of the edge e_f.

We now want to show that this map is surjective. For this, it suffices to show that every star class involving two moving particles is in the image of this map by the discussion above. Given a cycle representing such a star class with moving particles $\{s_1,s_2\}$ we can forget about the fixed particles. This cycle is then homologous to a linear combination of cycles in $C_1(\mathrm{Conf}_{\{s_1,s_2\}}(G))$ representing the chosen basis elements of $H[2]_{\{s_1,s_2\}}$. It is now sufficient to show that if we add for each of those cycles the fixed particles onto e_f, then the corresponding cycles in $C_1(\mathrm{Conf}_S(G))$ are still homologous. We will therefore show that whenever we have a chain bounding a linear combination of star cycles in $\mathrm{Conf}_{\{s_1,s_2\}}(G)$, we find a chain bounding the corresponding cycle in $\mathrm{Conf}_S(G)$ with the fixed particles $\underline{\mathbf{n}}-\{s_1,s_2\}$ added to e_f.

For easier description we subdivide the edge e_f by adding a vertex v_f of valence 2, and we denote the two parts of e_f by e_f^1 and e_f^2. The combinatorial model constructed in Proposition 1.20 works in the same way if not every vertex is essential, and we will use this 2-dimensional combinatorial model of $\mathrm{Conf}_{\{s_1,s_2\}}(G)$ to prove the claim. Let a linear combination $X=\Sigma\alpha_i X_i$ of cellular cycles representing star classes in this combinatorial model and a cellular chain Y bounding X be given. Remove the interior of all cells of Y where a particle moves from e_f^1 to v_f, giving a new chain $\tilde{Y}$. Since for no star class any particle moves towards v_f, this does not remove any cells in the boundary of Y. The new boundary has four parts, each of which is determined by the particle that is fixed and the position of this particle, which is either on v_f or on e_f^1. In each of these components the other particle moves along a 1-cycle in G, and each such cycle appears with different signs once with the fixed particle on v_f and once with the fixed particle on e_f^1.

Since in this modified chain $\tilde{Y}$ no particle moves from e_f^1 to v_f, this determines a well-defined chain in $C_1(\mathrm{Conf}_S(G))$ by putting the fixed particles $\underline{\mathbf{n}}-\{s_1,s_2\}$ onto e_f^1. This chain bounds X with the fixed particles added to e_f and the linear combination of circle classes with one of s_1 and s_2 fixed on v_f or e_f^1. Therefore, it is sufficient to show the following: given a 1-cycle in $C_1(\mathrm{Conf}_{\{s_1\}}(G))$ not meeting

e_f, the 1-cycle in $C_1(\mathrm{Conf}_S(G))$ given by the difference of putting the remaining particles onto e_f in two different orderings is null-homologous. Writing the 1-cycle as a linear combination of cycles represented by embedded circles, this follows from Proposition 4.28.

Therefore, the map (4.2) is surjective.

The composition of (4.2) with the sum of the projection maps to one and two particles

$$H_1(\mathrm{Conf}_S(G);A) \to \bigoplus_{s_1\in\underline{n}} H_1(\mathrm{Conf}_{\{s_1\}}(G);A) \oplus \bigoplus_{\{s_1,s_2\}\subset\underline{n}} H[2]_{\{s_1,s_2\}}$$

is by construction the identity, which shows that (4.2) is in fact an isomorphism. Therefore, $H_1(\mathrm{Conf}_S(G);A)$ is torsion-free.

It is straightforward to check that this is actually an isomorphism of $\mathbf{FI}^{\mathrm{op}}$-modules with the obvious $\mathbf{FI}^{\mathrm{op}}$-module structure on the source of (4.2): for each of the basic classes above all fixed particles are in the star of one vertex which is disjoint from the paths of the moving particles, so permuting fixed particles does not change the homology classes in $H_1(\mathrm{Conf}_S(G))$. Therefore, the map (4.2) is compatible with permutation of fixed particles, and since it is also compatible with inclusions this shows that this isomorphism is an isomorphism of $\mathbf{FI}^{\mathrm{op}}$-modules.

Applying $\mathrm{Hom}_A(\bullet,A)$ to this map shows that

$$H^1(\mathrm{Conf}_S(G);A) \cong \bigoplus_{s_1\in\underline{n}} H^1(\mathrm{Conf}_{\{s_1\}}(G);A) \oplus \bigoplus_{\{s_1,s_2\}\subset\underline{n}} \mathrm{Hom}_A(H[2]_{\{s_1,s_2\}},A).$$

Varying n, this gives an isomorphism of **FI**-modules. The right hand side is clearly finitely generated in degree 2, so the same is true for $H^1(\mathrm{Conf}_\bullet(G);A)$. □

In the proof, we in fact identified the first homology group explicitly:

Proposition 4.30. *For a graph G and a module A as in Theorem I, there exists an isomorphism of $A\Sigma_n$-modules*

$$H_1(\mathrm{Conf}_n(G);A) \cong \bigoplus_{s_1\in\underline{n}} H_1(\mathrm{Conf}_{\{s_1\}}(G);A) \oplus \bigoplus_{\{s_1,s_2\}\subset\underline{n}} H[2]_{\{s_1,s_2\}}$$

for modules $H[2]_{\{s_1,s_2\}} \subset H_1(\mathrm{Conf}_{\{s_1,s_2\}}(G);A)$ generated by star classes.

Bibliography

[Abr00] Aaron David Abrams. *Configuration Spaces and Braid Groups of Graphs*. PhD thesis, UC Berkley, 2000.

[AG02] Aaron Abrams and Robert Ghrist. Finding topology in a factory: configuration spaces. *Amer. Math. Monthly*, 109(2):140–150, 2002.

[BF09] Kathryn Barnett and Michael Farber. Topology of configuration space of two particles on a graph. I. *Algebr. Geom. Topol.*, 9(1):593–624, 2009.

[BH99] Martin R. Bridson and André Häfliger. *Metric Spaces of Non-Positive Curvature*. Springer, 1999.

[CEF15] Thomas Church, Jordan S. Ellenberg, and Benson Farb. FI-modules and stability for representations of symmetric groups. *Duke Math. J.*, 164(9):1833–1910, 2015.

[CF13] Thomas Church and Benson Farb. Representation theory and homological stability. *Adv. Math.*, 245:250–314, 2013.

[Che16] Safia Chettih. *Dancing in the Stars: Topology of Non-k-equal Configuration Spaces of Graphs*. PhD thesis, University of Oregon, 2016.

[CL16] Safia Chettih and Daniel Lütgehetmann. Configuration Spaces of Graphs, 2016, arXiv:1612.08290.

[Far06] Daniel Farley. Homology of tree braid groups. In *Topological and asymptotic aspects of group theory*, volume 394 of *Contemp. Math.*, pages 101–112. Amer. Math. Soc., Providence, RI, 2006.

[FH91] W. Fulton and J. Harris. *Representation Theory: A First Course*. Graduate Texts in Mathematics / Readings in Mathematics. Springer New York, 1991.

[FH10] Michael Farber and Elizabeth Hanbury. Topology of configuration space of two particles on a graph, II. *Algebr. Geom. Topol.*, 10(4):2203–2227, 2010.

[FS05] Daniel Farley and Lucas Sabalka. Discrete Morse theory and graph braid groups. *Algebr. Geom. Topol.*, 5:1075–1109, 2005.

[Ghr01] Robert Ghrist. Configuration spaces and braid groups on graphs in robotics. In *Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998)*, volume 24 of *AMS/IP Stud. Adv. Math.*, pages 29–40. Amer. Math. Soc., Providence, RI, 2001.

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge: Cambridge University Press, 2002.

[KKP12] Jee Hyoun Kim, Ki Hyoung Ko, and Hyo Won Park. Graph braid groups and right-angled Artin groups. *Trans. Amer. Math. Soc.*, 364(1):309–360, 2012.

[KP12] Ki Hyoung Ko and Hyo Won Park. Characteristics of graph braid groups. *Discrete Comput. Geom.*, 48(4):915–963, 2012.

[Lü14] Daniel Lütgehetmann. Configuration spaces of graphs. Master's thesis, Freie Universität Berlin, 2014.

[Lü17] Daniel Lütgehetmann. Representation Stability for Configuration Spaces of Graphs, 2017, arXiv:1701.03490.

[McC01] John McCleary. *A user's guide to spectral sequences. 2nd ed.* Cambridge: Cambridge University Press, 2nd ed. edition, 2001.

[MS17] Tomasz Maciążek and Adam Sawicki. Homology groups for particles on one-connected graphs. *J. Math. Phys.*, 58(6):062103, 24, 2017.

[Ram16a] Eric Ramos. Generalized representation stability and FI_d-modules, 2016, arXiv:1606.02673.

[Ram16b] Eric Ramos. Stability phenomena in the homology of tree braid groups, 2016, arXiv:1609.05611.

[Ram17] Eric Ramos. Configuration spaces of graphs with certain permitted collisions, March 2017, arXiv:1703.05535.

[Św01] Jacek Świątkowski. Estimates for homological dimension of configuration spaces of graphs. *Colloquium Mathematicum*, 89(1):69–79, 2001.

Abstract

For a finite graph G and a natural number n we study the homology of the configuration space $\mathrm{Conf}_n(G)$ of n particles in G. A graph is called a "tree with loops" if it can be constructed from a tree by taking the iterated wedge sum with copies of S^1 for different choices of base points. We prove that if G is a tree with loops then the homology of $\mathrm{Conf}_n(G)$ is torsion-free and generated by products of 1-dimensional classes. For general graphs G we give a generating set for the first homology group $H_1(\mathrm{Conf}_n(G))$. Using these results and the techniques used in their proofs we then prove representation stability for specific sequences of configuration spaces of graphs given by either enlarging the graph or increasing the number of particles.

Zusammenfassung

Für einen endlichen Graphen G und eine natürliche Zahl n untersuchen wir die Homologie des Konfigurationsraums $\mathrm{Conf}_n(G)$ von n Partikeln in G. Wir nennen einen endlichen Graphen einen "Baum mit Schleifen" wenn er durch Ankleben (per Wedge-Produkt an verschiedenen Basispunkten) von Kopien des Kreises S^1 an einen Baum konstruiert werden kann. Wir beweisen, dass für einen Baum mit Schleifen G die Homologie von $\mathrm{Conf}_n(G)$ torsionsfrei und erzeugt von Produkten 1-dimensionaler Homologieklassen ist. Für allgemeine Graphen G geben wir ein Erzeugendensystem für die erste Homologiegruppe $H_1(\mathrm{Conf}_n(G))$ an. Mit diesen Resultaten und den in den Beweisen benutzten Techniken zeigen wir anschließend Darstellungsstabilität für bestimmte Folgen von Konfigurationsräumen von Graphen, welche entweder durch Vergrößerung des Graphen oder Erhöhung der Partikelanzahl definiert werden.

Lebenslauf

Der Lebenslauf wird aus Gründen des Datenschutzes in der elektronischen Fassung meiner Arbeit nicht veröffentlicht.

Selbstständigkeitserklärung

Hiermit erkläre ich, dass ich die vorliegende Arbeit selbstständig angefertigt und ausschließlich die angegebenen Quellen und Hilfsmittel verwendet habe. Außerdem versichere ich, dass die Arbeit nicht schon einmal in einem früheren Promotionsverfahren eingereicht worden ist.

Berlin, 29. August 2017 Daniel Lütgehetmann